Facteurs cycliques, automorphismes irréductibles et axes contractants pour les automorphismes de groupes de Baumslag-Solitar généralisés

Chloé Papin

UNIVERSITÉ DE RENNES 1

THÈSE DE DOCTORAT DE

L'UNIVERSITÉ DE RENNES 1

ÉCOLE DOCTORALE N° 601
Mathématiques et Sciences et Technologies de l'Information et de la Communication
Spécialité : *Mathématiques et Interactions*

Par
Chloé PAPIN

Facteurs cycliques, automorphismes irréductibles et axes contractants pour les groupes de Baumslag-Solitar généralisés

Thèse présentée et soutenue à Rennes, le 10 janvier 2022
Unité de recherche : IRMAR (UMR CNRS 6625)

Rapporteurs avant soutenance :

Yael ALGOM-KFIR Senior Lecturer, University of Haifa
Ilya KAPOVICH Professor, Hunter College of CUNY

Composition du Jury :

Président :	Yves CORNULIER	Directeur de Recherche CNRS, Université de Lyon 1
Examinateurs :	Yael ALGOM-KFIR	Senior Lecturer, University of Haifa
	Yves CORNULIER	Directeur de Recherche CNRS, Université de Lyon 1
	Rémi COULON	CNRS, Université de Rennes 1
	Cornelia DRUTU	Professor, University of Oxford
Dir. de thèse :	Vincent GUIRARDEL	Professeur, Université de Rennes 1

Remerciements

Je tiens à remercier en tout premier lieu mon directeur de thèse, Vincent Guirardel, pour m'avoir encadrée durant ces trois années de thèse. Merci pour les nombreuses discussions qui m'ont permis de faire avancer ce travail, pour la rigueur des relectures et pour les suggestions qui, souvent, apportaient un éclairage nouveau sur les problèmes que nous rencontrions.

Merci à Yael Algom-Kfir et Ilya Kapovich d'avoir accepté d'écrire un rapport sur ma thèse. Je remercie aussi Yves Cornulier, Rémi Coulon et Cornelia Drutu de s'être intéressés à mes travaux et d'avoir accepté de faire partie de mon jury.

Je ne serais pas arrivée en thèse sans Camille Horbez, qui a encadré mon stage de master et m'a mise en relation avec Vincent. Pour cela je lui suis reconnaissante.

J'ai beaucoup apprécié le temps que j'ai passé à l'IRMAR, en grande partie grâce à mes collègues, qu'il s'agisse des chercheurs et chercheuses ou bien du personnel administratif. J'ai passé d'excellents moments avec eux. Je n'oublie pas Élodie Cottrel à l'école doctorale, qui m'a bien aidée dans mes démarches administratives.

Merci à mes amis et collègues doctorants Lucien, Titouan et Victor pour nos discussions mathématiques ou non, et pour leur soutien important pendant ces deux dernières années. Merci aussi à Josselin d'être souvent venu nous offrir l'occasion de plier des icosaèdres triambiques et des éléphants. Sans oublier Thi et Vincent, qui m'ont empêchée de mourir de faim alors que je rédigeais, ni Simon... J'oublie beaucoup de monde ; je n'en suis pas moins reconnaissante envers les doctorants qui m'ont accompagnée au long de ma thèse.

Enfin je voudrais remercier ma famille pour leurs encouragements et leur soutien pendant ma thèse mais aussi tout au long de mes études. Merci à Léo pour son aide et ses conseils lorsque j'étais en difficulté. Merci aussi à ces quelques personnes que je retrouve de temps en temps autour d'une tisane réglisse-menthe, dont plusieurs écrivent comme moi des thèses, et à mes amis Guillaume et Fabien qui m'ont précédée dans la voie du doctorat.

TABLE DES MATIÈRES

INTRODUCTION

Le groupe de Baumslag-Solitar $\mathrm{BS}(p,q)$, où p,q sont deux entiers non nuls, est le groupe défini par la présentation suivante :

$$\mathrm{BS}(p,q) := \langle a,t | ta^pt^{-1} = a^q \rangle.$$

C'est une extension HNN de $\mathbb{Z}$ au-dessus $\mathbb{Z}$. Selon le choix des paramètres p et q, les groupes obtenus peuvent avoir des propriétés variées. Ils ont été introduits par Baumslag et Solitar dans [BS62] en tant qu'exemples de groupes non hopfiens, c'est-à-dire possédant des endomorphismes surjectifs non injectifs. Par exemple $\mathrm{BS}(2,3)$ n'est pas hopfien, mais $\mathrm{BS}(2,4)$ l'est, bien qu'il ne soit pas résiduellement fini ([CL83]).

On appelle *groupes de Baumslag-Solitar généralisés* (GBS) les groupes isomorphes au groupe fondamental d'un graphe de groupes fini dont tous les groupes de sommets et d'arêtes sont cycliques infinis. La notion de graphes de groupes, développée par Serre dans [Ser77], est une manière d'encoder les groupes fondamentaux de certains espaces : la définition de leur groupe fondamental est une variante du groupe fondamental usuel.

Les groupes de Baumslag-Solitar appartiennent à cette famille, $\mathrm{BS}(p,q)$ étant le groupe fondamental du graphe de la figure 1. La même figure présente également l'*arbre de Bass-Serre* associé ou *revêtement universel du graphe de groupes*, qui est l'analogue du revêtement universel usuel : il admet une action du groupe fondamental du graphe de groupes, mais au lieu d'être libre, cette action admet des stabilisateurs de sommets ou d'arêtes qui sont donnés à conjugaison près par les étiquettes du graphe de groupes.

Prenons l'exemple de $\mathrm{BS}(p,q)$, illustré par la figure 2. On prend un cercle S^1 et un anneau $S^1 \times [0,1]$. On recolle l'un des deux bords de l'anneau sur le cercle en faisant p tours. On recolle l'autre bord en faisant q tours avec la même orientation. L'espace obtenu a pour groupe fondamental $\mathrm{BS}(p,q)$.

Pour un graphe de groupes cycliques plus compliqué Γ, on prend un cercle par sommet et un anneau par arête. On recolle de la même manière les anneaux sur les cercles avec un nombre de tours non nul correspondant à l'indice de l'inclusion du groupe d'arête dans le groupe de sommet. On obtient ainsi un espace dont le groupe fondamental est $\pi_1(\Gamma)$. Un

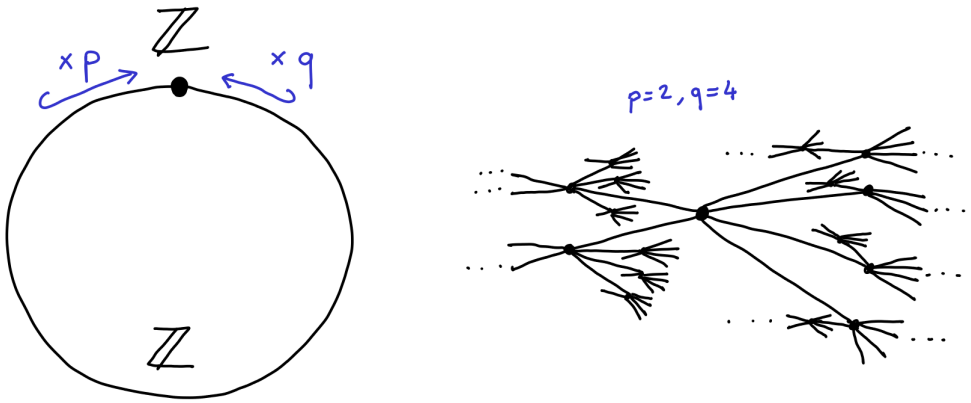

FIGURE 1 – Un graphe de groupes dont $\mathrm{BS}(p,q)$ est le groupe fondamental, arbre de Bass-Serre associé.

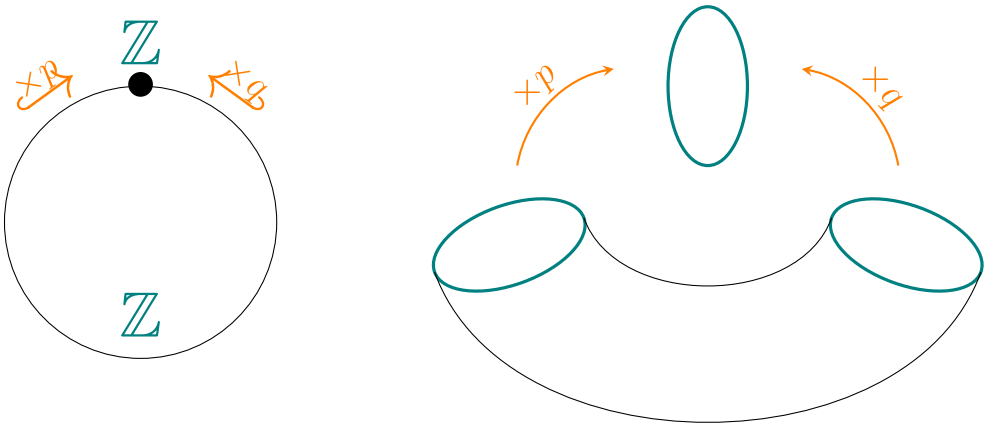

FIGURE 2 – Recollement d'anneaux et de cercles et graphe de groupes associé

groupe de Baumslag-Solitar généralisé est le groupe fondamental d'un espace construit par un tel recollement d'anneaux sur des cercles.

Dans les graphes de groupes à groupes de sommets et d'arêtes cycliques infinis, on choisit un générateur pour chaque groupe de sommet et d'arête. L'inclusion d'un groupe d'arête dans un groupe de sommet est donnée par un entier signé non nul. Ainsi on peut encoder ces graphes de groupes par des *graphes étiquetés*, qui sont des graphes dans lesquels chaque arête est munie de deux étiquettes, une à chaque extrémité, portant chacune un entier non nul représentant l'inclusion.

L'écriture d'un groupe GBS comme groupe fondamental d'un graphe de groupes n'est pas unique en général : une infinité de graphes de groupes différents peuvent avoir des

groupes fondamentaux isomorphes. Le problème d'isomorphisme consiste à décider si deux graphes étiquetés Γ_1, Γ_2 ont des groupes fondamentaux isomorphes. On ne connaît pas de solution générale à ce problème. Il existe des solutions dans certains cas particuliers : si le premier nombre de Betti de Γ_1 et Γ_2 est 1, Clay et Forester prouvent dans [CF08a] qu'on peut décider si $\pi_1(\Gamma_1) \simeq \pi_1(\Gamma_2)$. Dans [Lev07] Levitt définit un critère sur les graphes étiquetés qui permet d'identifier les groupes GBS dont le groupe d'automorphismes ne contient pas F_2, et est donc, dans un sens, assez petit. Pour ces groupes GBS, il montre que le problème d'isomorphisme a une solution.

D'après la théorie de Bass-Serre, il existe une équivalence entre

— les groupes fondamentaux de graphes de groupes dont tous les groupes de sommets et d'arêtes sont $\mathbb{Z}$

— les groupes qui agissent sur un arbre simplicial avec stabilisateurs de sommets et d'arêtes isomorphes à $\mathbb{Z}$

Ici, plus particulièrement, on s'intéresse au groupe des automorphismes extérieurs des groupes de Baumslag-Solitar (généralisés). Le groupe des automorphismes extérieurs d'un groupe G est

$$\mathrm{Out}(G) := \mathrm{Aut}(G)/\mathrm{Inn}(G)$$

où $\mathrm{Aut}(G)$ est le groupe des automorphismes de G et $\mathrm{Inn}(G) := \{c_g : x \mapsto gxg^{-1}, g \in G\}$ est le sous-groupe distingué des automorphismes intérieurs.

Selon G, les groupes d'automorphismes extérieurs peuvent être très différents. Par exemple, $\mathrm{Out}(\mathrm{BS}(2,3))$ est fini. Au contraire $\mathrm{Out}(\mathrm{BS}(2,4))$ n'est pas de type fini et contient un groupe libre non abélien. Collins donne une présentation finie de $\mathrm{Aut}(\mathrm{BS}(p,q))$ lorsque p et q sont premiers entre eux dans [Col78]. Dans [Gil+00], Gilbert, Howie, Metaftsis et Raptis en donnent une présentation lorsque p et q ne sont pas multiples l'un de l'autre ; le groupe $\mathrm{Out}(\mathrm{BS}(p,q))$ est alors dihédral d'ordre $2|p-q|$. Dès que q est un multiple de p et $|p| > 1$ en revanche, le groupe $\mathrm{Out}(\mathrm{BS}(p,q))$ n'est même pas de type fini. Collins et Levin ont montré ce résultat de manière algébrique et trouvé une présentation du groupe d'automorphismes dans [CL83]. Plus tard, Clay dans [Cla09] a également montré ce résultat en utilisant des méthodes géométriques, en utilisant un *espace de déformation* que nous définirons plus tard.

On peut s'attendre à ce que les groupes d'automorphismes des groupes GBS soient également très variés. Dans cette thèse l'approche utilisée pour les étudier est une approche géométrique. À l'instar des groupes libres, les groupes GBS ont beaucoup d'actions sur

des arbres, et notre étude de $\mathrm{Out}(G)$ sera basée sur un espace dont les points sont des actions de G sur des arbres avec certaines restriction sur les stabilisateurs de sommets et d'arêtes. C'est ce qu'on appelle l'espace de déformation. C'est un espace contractile muni d'une action naturelle de $\mathrm{Out}(G)$.

Enfin, un moyen d'en apprendre davantage sur $\mathrm{Out}(G)$ est de trouver une action intéressante de $\mathrm{Out}(G)$ sur un espace hyperbolique au sens de Gromov. L'analogie avec l'étude de $\mathrm{Out}(F_N)$ est une motivation. En effet, $\mathrm{Out}(F_N)$ admet une action sur le complexe des facteurs libres, défini dans [HV98] et dont l'hyperbolicité a été prouvée dans [BF14]. Dans cette action, les automorphismes qui agissent de façon hyperbolique sont les automorphismes *complètement irréductibles*. Il existe par ailleurs d'autres complexes hyperboliques intéressants sur lesquels $\mathrm{Out}(F_N)$ agit. Cette action de $\mathrm{Out}(F_N)$ est également à rapprocher de l'action du groupe modulaire d'une surface $\mathrm{Mod}(S_g)$ sur le complexe des courbes, qui en est l'inspiration. L'hyperbolicité du complexe des courbes a été montrée dans [MM99], et les éléments qui agissent dessus de façon loxodromique sont les éléments dits *pseudo-Anosov* de $\mathrm{Mod}(S_g)$.

Espace de déformation, un analogue de l'outre-espace

Un groupe GBS G agit sur des arbres et il existe généralement beaucoup d'actions de G sur des arbres. L'espace de déformation est un espace dont les points correspondent à certaines de ces actions.

Une action sur un arbre peut se représenter par un graphe de groupes, qui a l'avantage d'être fini. D'après la théorie de Bass-Serre ([Ser77]), le groupe $\pi_1(\Gamma)$ où Γ est un graphe de groupes agit sur un arbre (appelé le *revêtement universel* de Γ) avec des stabilisateurs de sommets et d'arêtes cycliques infinis. Réciproquement, la donnée de l'action d'un groupe G sur un arbre avec des stabilisateurs cycliques infinis permet de retrouver le graphe de groupes, dont le groupe fondamental est isomorphe à G. Un graphe de groupes *marqué*, c'est-à-dire un graphe Γ muni d'une identification entre $\pi_1(\Gamma)$ et G, ou un graphe étiqueté marqué, contient la même information que l'action de G sur le revêtement universel de Γ.

Partant d'une action ρ de G sur un arbre T, il existe plusieurs moyens de la modifier pour en faire une nouvelle action. On peut tordre l'action ρ par un automorphisme $\phi \in \mathrm{Aut}(G)$ pour obtenir une nouvelle action $\rho_\phi = \rho \circ \phi$ sur T. En outre, il existe en général une infinité de graphes de groupes ayant un groupe fondamental isomorphe à G.

Plus généralement, on peut effectuer des *pliages* (de Stallings) pour construire de

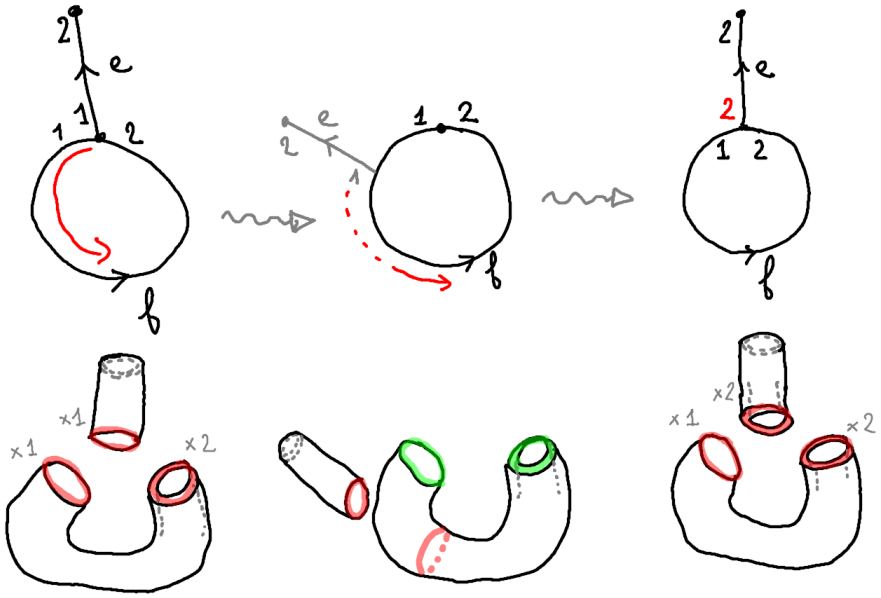

FIGURE 3 – Glissement de l'arête e sur l'arête f et illustration en termes de cercles et d'anneaux. Les cercles de même couleur sont recollés ensemble.

nouvelles actions. Donnons un exemple d'une telle construction.

Exemples A. Soit $G := \mathrm{BS}(2,4) = \langle a, t | ta^2t^{-1} = a^4 \rangle$. Le graphe en haut à gauche de la figure 3 représente le groupe G (bien qu'il soit différent du graphe de la figure 1). On peut faire *glisser* l'arête e le long de la boucle f, c'est-à-dire qu'on déplace son point d'attache le long de f. Le dessin du bas représente l'opération dans l'espace fait d'anneaux et cercles : on fait glisser le cercle d'attache de l'anneau du haut le long de l'anneau de la boucle (en rouge). Comme le nombre de tours à chacun des bouts de l'anneau est différent, après le glissement, le cercle d'attache fait deux fois le tour du cercle d'attache initial.

Les espaces de déformations ont été introduits par Forester ([For06]). Ils regroupent l'ensemble des actions de G sur des arbres avec certains stabilisateurs. Pour des groupes GBS, on peut les définir de la façon suivante.

Définition B. [For02, Corollaire 6.10] Soit G un groupe GBS non élémentaire, c'est-à-dire non isomorphe à $\mathbb{Z}$, $\mathbb{Z}^2$ ou le groupe fondamental d'une bouteille de Klein. Soit $\mathcal{T}$ l'ensemble des actions minimales par isométries de G sur des arbres métriques simpliciaux T à stabilisateurs de points cycliques infinis. Soit $\sim$ la relation d'équivalence sur $\mathcal{T}$ minimale telle que $T \sim T'$ s'il existe une homothétie G-équivariante qui envoie T sur T'. On appelle *espace de déformation (cyclique)* l'ensemble $\mathcal{D} := \mathcal{T}/\sim$.

On munit $\mathcal{D}$ d'une topologie. L'espace des actions minimales de G sur des arbres métriques simpliciaux a une structure naturelle d'espace simplicial, à laquelle est associée la *topologie faible*. Comme $\mathcal{D}$ s'identifie à la réunion d'une partie des simplexes, il hérite de la topologie faible. L'espace $\mathcal{D}$ hérite également de la *topologie des fonctions longueur* (ou *topologie des axes*) : pour un arbre T, une fonction longueur est $l_T : G \to \mathbb{R}, g \mapsto \|g\|_T$ [MS84]. Cette topologie est équivalente à la *topologie de Gromov équivariante* [Pau89]. La topologie faible et la topologie des axes sont équivalentes sur $\mathcal{D}$, ce qui est une conséquence du fait que les arbres de $\mathcal{D}$ sont localement finis [GL07, Proposition 5.4].

L'espace $\mathcal{D}$ admet une action de $\mathrm{Aut}(G)$ par précomposition de l'action telle que $\mathrm{Inn}(G)$ agit trivialement sur $\mathcal{D}$. Ainsi la précomposition définit bien une action de $\mathrm{Out}(G)$ sur $\mathcal{D}$.

On peut considérer $\mathcal{D}$ comme un analogue de l'outre-espace CV_N pour le groupe d'automorphismes du groupe libre $\mathrm{Out}(F_N)$, ou encore comme un analogue de l'espace de Teichmüller pour un groupe modulaire de surface.

L'espace $\mathcal{D}$ est essentiellement un complexe simplicial contractile ([GL07]).

Automorphismes irréductibles, facteurs cycliques

Une question qui s'est révélée cruciale pour l'étude des automorphismes de groupes libres est l'existence de facteurs libres préservés par un automorphisme.

Définition C. Soit F un groupe. On dit qu'un sous-groupe $A \leq F$ est un *facteur libre* de F s'il existe un sous-groupe B tel que F est le produit libre $A * B$.

Remarque D. Prenons un isomorphisme $F_N \simeq \pi_1(\Gamma)$ où Γ est un graphe. Si Γ' est un sous-graphe connexe non vide de Γ, on a une inclusion $\pi_1(\Gamma') \hookrightarrow \pi_1(\Gamma)$ qui dépend d'un choix de point base (voir la figure 4). Par le théorème de Van Kampen, on peut montrer que le sous-groupe de F_N correspondant est un facteur libre. Réciproquement, tout facteur libre peut être obtenu de cette manière.

Il existe une manière de voir les facteurs libres au revêtement universel. Lorsqu'on relève Γ' dans le revêtement universel de Γ, qui est un arbre, on obtient un sous-arbre. Tous les translatés de ce sous-arbre sous l'action de $\pi_1(\Gamma)$ sont disjoints deux à deux. On peut les *écraser*, c'est-à-dire remplacer chaque sous-arbre par un sommet sur lequel on rattache les arêtes. On obtient une nouvelle action de F_N sur un arbre et le stabilisateur de chacun des nouveaux sommets créés est un conjugué de $\pi_1(\Gamma')$. Réciproquement, le stabilisateur d'un sous-arbre disjoint de tous ses translatés est un facteur libre.

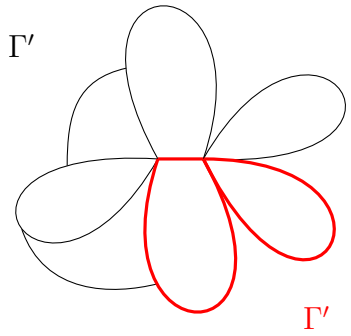

FIGURE 4 – Un facteur libre vu comme sous-graphe d'un graphe de groupe fondamental F_N

Définition E. Soit $\phi \in \mathrm{Out}(F_N)$. On dit que ϕ est *complètement irréductible* si aucune puissance de ϕ ne préserve la classe de conjugaison d'un facteur libre différent de $\{1\}$ ou F_N.

Les automorphismes complètement irréductibles sont un analogue des éléments pseudo-Anosov pour les groupes modulaires. Ils agissent de façon loxodromique sur l'outre-espace et possèdent ce qu'on appelle des *représentants train track*, qui constituent un axe de translation. Les représentants train track sont un outil indispensable à l'étude de ces automorphismes.

Dans le cas des groupes GBS, les facteurs libres ne sont pas la bonne notion. Une notion analogue qui convient dans ce cas est la notion de *facteur cyclique* (voir définition 1.2.2).

Définition F. Soit G un GBS. On dit que A est un *facteur cyclique* de G s'il existe un graphe de groupes cycliques Γ avec une identification $G \simeq \pi_1(\Gamma)$ (au sens des graphes de groupes) et un sous-graphe connexe $\Gamma' \subset \Gamma$ tel que A s'identifie au sous-groupe $\pi_1(\Gamma') \subset \pi_1(\Gamma)$ (défini à conjugaison près).

Remarques G.
- — De façon équivalente, A est un *facteur cyclique* de G s'il existe un arbre $T \in \mathcal{D}$ et un écrasement $T \to T'$ tel que A est un stabilisateur de sommet de T'.
- — On peut définir un raffinement de la notion de facteur cyclique, qui consiste à imposer une condition supplémentaire sur les groupes d'arêtes du graphe de groupes

Γ. Nous appelons cela un *facteur spécial*. Un facteur spécial est toujours un facteur cyclique.

On peut alors définir des automorphismes complètement irréductibles en remplaçant « facteur libre »par « facteur cyclique »(voire « facteur spécial », ce qui donne une notion un peu différente).

Lorsqu'on étudie un automorphisme $\phi \in \mathrm{Out}(G)$, on peut considérer un *représentant* qui est une application $f : T \to T$ vérifiant $f(gt) = \phi(g)f(t)$ pour tout $t \in T$ et tout $g \in G$. Il peut arriver que ce représentant laisse invariante une sous-forêt G-invariante propre $F \subset T$. En écrasant cette forêt, c'est-à-dire en identifiant les points de T qui appartiennent à une même composante connexe de la forêt, on obtient un nouvel arbre. Si la forêt possède une composante non bornée alors le stabilisateur de cette composante est un facteur cyclique et c'est un stabilisateur de sommet dans l'arbre écrasé. Sa classe de conjugaison est alors périodique par l'automorphisme : l'automorphisme est alors réductible.

Des représentants train track existent dans certains cas particuliers. Pour les groupes de Baumslag-Solitar par exemple, Bouette montre dans [Bou16] que tout automorphisme est réductible. En revanche, quitte à changer d'espace de déformation en autorisant des stabilisateurs de sommets plus gros, il existe des automorphismes irréductibles et ceux-ci admettent des train tracks dans le nouvel espace de déformation. Un autre exemple est celui des groupes GBS sans module entier non trivial. Pour ceux-ci, il existe une borne sur le nombre d'orbites d'arêtes pour tous les éléments de $\mathcal{D}$ ([For06]), ce qui permet de prouver l'existence de train tracks pour tout automorphisme irréductible. L'existence d'automorphismes irréductibles en revanche n'est pas systématique !

Dans le cas général, on ne sait pas si l'existence de représentants train track pour les automorphismes complètement irréductibles est toujours vraie.

Comprendre les facteurs cycliques

Dans le cas des groupes libres, il existe un algorithme qui permet de décider, étant donné un élément $g \in F_N$, s'il existe un facteur libre qui le contient, et le cas échéant, de trouver un tel facteur libre.

La question se reformule en termes de facteurs spéciaux pour les groupes de Baumslag-Solitar et nous y répondons positivement :

Théorème H. *Il existe un algorithme qui prend en entrée un groupe GBS sous forme de graphe de groupes et un élément loxodromique $g \in G$ sous forme d'un chemin dans*

le graphe de groupes, et qui retourne soit une preuve qu'aucun facteur cyclique propre ne contient g, soit un facteur cyclique propre qui contient g sous forme d'un sous-graphe du graphe de groupes.

Un élément contenu dans un facteur cyclique propre est appelé un élément *simple*.

Nous montrons en fait une version plus forte de ce résultat (théorème 1.3.3). Un *système de facteurs cycliques* se définit comme un sous-graphe éventuellement non connexe d'un graphe de groupes Γ ; c'est l'ensemble des sous-groupes de G qui s'identifient au groupe fondamental d'une des composantes connexes. On peut y penser comme un ensemble de facteurs cycliques « compatibles entre eux ». Le théorème ci-dessus peut s'appliquer à une famille finie $g_1, \ldots, g_n$ d'éléments loxodromiques de G, et permet alors de décider s'il existe un système de facteurs cycliques tel que chaque g_i est contenu dans un membre du système de facteurs cycliques.

Notons que le graphe de groupes retourné par l'algorithme n'est en général pas un sous-graphe du graphe de groupes initial. En effet, il existe en général une infinité de définitions de G comme groupe fondamental d'un graphe de groupes et il n'est pas évident de savoir si l'on peut en trouver une dans laquelle les g_i sont dans un sous-graphe. Le principe de l'algorithme est de trouver un graphe de groupes particulier pour G dans lequel on peut voir le facteur cyclique recherché (ou le système de facteurs cycliques) comme sous-graphe, ou bien de montrer qu'un tel graphe de groupes n'existe pas.

Les motivations pour un tel algorithme sont diverses. On peut rapprocher la question ci-dessus du problème de conjugaison par exemple : étant donnés g et h dans G, existe-t-il $a \in G$ tel que $g = aha^{-1}$? Ce dernier problème a été résolu par Beeker dans [Bee15]. On peut également la rapprocher du problème de Whitehead consistant à décider si deux éléments du groupe sont dans la même orbite sous l'action du groupe d'automorphismes : étant donnés g et h, existe-t-il $\phi \in \mathrm{Aut}(G)$ tel que $\phi(g) = h$? Ce problème est actuellement ouvert.

Le théorème s'appuie sur un critère très facile à mettre en œuvre pour garantir qu'un élément $g \in G$ n'est pas simple, ce que nous utiliserons plus tard.

D'autre part, les méthodes utilisées peuvent s'adapter dans une certaine mesure à d'autres contextes. Pour étudier les automorphismes irréductibles, nous utiliserons une variante de cet algorithme concernant les laminations, par exemple.

La méthode utilisée pour l'algorithme est basée sur les *graphes de Whitehead*. Étant donné $g \in G$ loxodromique, un arbre $T \in \mathcal{D}$ et un sommet $v \in T$, le graphe de Whitehead

$\mathrm{Wh}_T(g, v)$ est le graphe défini ainsi :

— les sommets sont les arêtes de T avec sommet initial v

— on relie e et e' s'il existe un conjugué de g dont l'axe contient e et e'.

Les graphes de Whitehead sont une manière d'encoder le trajet des axes des conjugués de g dans l'arbre. Ils permettent de savoir quelles modifications de l'arbre on peut effectuer sans changer la longueur de translation de g. L'algorithme est basé sur des *dépliages* et *éclatements*, qui sont des modifications de l'arbre qui augmentent son covolume mais préservent la longueur de translation de g ; la description de l'algorithme figure à la fin de la section 1.3 du chapitre 1.

Un facteur cyclique d'un groupe GBS est lui-même un groupe GBS. L'algorithme du théorème H peut être itéré pour trouver une suite de facteurs cycliques de plus en plus petits qui contiennent un élément g. Une telle suite est en fait nécessairement stationnaire. On a de plus :

Théorème I. *Soit $g \in G$ un élément loxodromique. Il existe un unique facteur cyclique minimal qui contient g.*

Si $g_1, \ldots, g_n$ est une collection finie d'éléments loxodromiques de G alors il existe un unique système de facteurs cycliques minimal qui contient $g_1, \ldots, g_n$.

On en déduit le corollaire suivant :

Corollaire J. *Il existe un algorithme qui prend en entrée un groupe GBS G sous forme de graphe de groupes et une collection finie $\mathcal{G}$ d'éléments loxodromiques de G, et qui trouve le plus petit système de facteurs cycliques contenant $\mathcal{G}$ sous la forme d'un sous-graphe d'un graphe de groupes pour G.*

Détection des automorphismes complètement irréductibles

La deuxième partie de ce travail est consacrée au problème suivant :

Problème K. *Existe-t-il un algorithme qui, étant donné un groupe GBS G et un automorphisme $\phi \in \mathrm{Out}(G)$ (ou $\phi \in \mathrm{Aut}(G)$), permet de décider si ϕ est complètement irréductible ?*

L'analogue pour les groupes libres existe. Kapovich en donne un dans [Kap14]. L'algorithme repose sur la recherche d'éléments périodiques pour l'automorphisme ainsi que sur le calcul des graphes de Whitehead de la lamination stable dans un représentant train track.

Un *représentant* de l'automorphisme $\phi \in \mathrm{Aut}(G)$ est une application $f : T \to T$ telle que pour tous $g \in G, t \in T$ on ait $f(gt) = \phi(g)f(t)$. Un *représentant train track* est un représentant tel que pour toute arête $e \in T$, pour tout $n \in \mathbb{N}$ l'image $f^n(e)$ soit une géodésique, c'est-à-dire ne présente pas de « demi-tour ». Les représentants train track ont été introduits par Bestvina et Handel dans [BH92]. Dans [FM15] Francaviglia et Martino étendent la notion aux automorphismes de produits libres et montrent l'existence de train tracks pour les automorphismes irréductibles. Dans [Mei15], Meinert montre également l'existence de train tracks pour le cas des groupes dont l'espace de déformation contient un nombre fini d'orbites de simplexes. Cela s'applique à quelques groupes GBS, mais pas à tous. Pour les groupes de Baumslag-Solitar $\mathrm{BS}(p, pn)$ Bouette montre l'existence de représentants train track dans un espace de déformation autre que l'espace de déformation cyclique. Dans ce cas, il n'y a pas d'automorphisme irréductible dans l'espace de déformation cyclique.

Les représentants train track sont des représentants privilégiés. D'une part, la propriété sur les itérés des arêtes est très puissante. D'autre part, ce sont des représentants qui minimisent $d_{\mathrm{Lip}}(T, T \cdot \phi)$ où d_{Lip} est la *métrique de Lipschitz* sur $\mathcal{D}$, que nous introduirons un peu plus loin.

La *lamination stable* est une famille de géodésiques bi-infinies de l'arbre obtenue en itérant f sur les arêtes, ce qui produit des segments de plus en plus longs et jamais « pliés »par f. Elle est stable par f.

Un résultat clé est le suivant :

Théorème L (Kapovich)**.** *Soit $\phi \in \mathrm{Out}(F_N)$ admettant un représentant train track irréductible $f : T \to T$ sans élément périodique simple. Alors ϕ est complètement irréductible si et seulement si tous les graphes de Whitehead de la lamination stable dans T sont connexes.*

Notons que si ϕ n'admet pas de représentant train track irréductible ou bien si ϕ possède un élément périodique simple, ϕ est réductible.

On a un énoncé analogue pour les groupes GBS, avec une nuance. La notion d'élément périodique est remplacée par celle d'élément *pseudo-périodique*, que nous détaillerons juste après :

Théorème M. *Soit G un groupe GBS. Soit $\phi \in G$ admettant un représentant train track irréductible $f : T \to T$ sans élément pseudo-périodique simple. Alors ϕ est complètement irréductible si et seulement si tous les graphes de Whitehead de la lamination stable dans T sont connexes.*

Un élément $g \in G$ est *pseudo-périodique* si $\|\phi^n(g)\|_T$ est borné pour $n \in \mathbb{N}$. Il existe une caractérisation des éléments pseudo-périodiques en termes de *chemins de Nielsen* qui sont une forme de chemins périodiques dans le représentant train track. En général, la classe de conjugaison d'un élément pseudo-périodique n'est pas périodique, mais la classe du facteur cyclique minimal le contenant est périodique (voir la proposition 2.3.11). C'est pourquoi l'existence d'éléments pseudo-périodiques simples implique la réductibilité de ϕ.

Alors que l'algorithme de Kapovich implique d'utiliser un algorithme pour calculer un représentant train track, par exemple celui de [BH92], on ne peut pas faire cela pour les groupes GBS car l'existence d'un représentant train track n'est pas connue.

On peut toutefois noter que les étapes de l'algorithme de Bestvina et Handel se transposent sans problème aux groupes GBS, bien que ce ne soit pas le cas des arguments qui permettent de prouver que l'algorithme termine. S'il termine toutefois, il donne bien un train track. Notre algorithme ne fonctionne pas sans la donnée d'un train track.

Pour appliquer le théorème M, il faut savoir faire plusieurs choses. Il faut pouvoir tester le fait que f est un représentant train track irréductible, ce qui est algorithmique. Plus complexe, il faut être capable de décider s'il existe des éléments pseudo-périodiques simples. On verra qu'on sait résoudre les deux problèmes suivants :

— Trouver les sous-groupes d'éléments pseudo-périodiques (proposition 2.4.4)

— Décider si un élément donné est simple (théorème 1.3.3)

mais on ne sait pas décider si un sous-groupe contient un élément simple. Enfin, il faut être capable de calculer les graphes de Whitehead de la lamination stable de manière algorithmique, ce qui est possible car l'arbre possède un nombre fini d'orbites de tournants et la lamination stable est quasi-périodique (lemme 2.1.21).

Ainsi on a le résultat suivant :

Théorème N. *Il existe un algorithme qui prend en entrée un groupe GBS G, un automorphisme $\phi \in \mathrm{Aut}(G)$ admettant un représentant train track $f : T \to T$, qui décide si ϕ admet des éléments dont la classe de conjugaison est pseudo-périodiques, et dans le cas contraire, il décide si la classe de ϕ dans $\mathrm{Out}(G)$ est complètement irréductible.*

Alors que les automorphismes de groupes de Baumslag-Solitar $\mathrm{BS}(p, pn)$ admettent des représentants train track dans un espace de déformation différent de l'espace de déformation cyclique (voir [Bou16]), il n'est malheureusement pas possible de leur appliquer l'algorithme ci-dessus. En effet, les démonstrations utilisent à plusieurs reprises le fait que les arbres de l'espace de déformation soient localement finis, ce qui n'est pas le cas dans l'espace de déformation modifié pour $\mathrm{BS}(p, pn)$.

Axes fortement contractants pour les automorphismes de groupes GBS

La dernière partie de cette thèse vise à montrer que $\mathcal{D}$ admet une certaine propriété de courbure négative, à savoir que les axes de certains automorphismes complètement irréductibles sont fortement contractants. Une telle propriété est montrée par Algom-Kfir dans [Alg11] pour l'outre-espace.

Commençons par définir sur $\mathcal{D}$ la *métrique de Lipschitz*. Soit $T, T' \in \mathcal{D}$ et $f : T \to T'$ une application lipschitzienne G-équivariante. On note $\mathrm{Lip}(f)$ la constante de Lipschitz de f.

On pose $\mathrm{Lip}(T, T') = \inf_{f:T\to T'} \mathrm{Lip}(f)$. On définit la métrique de Lipschitz par

$$d_{\mathrm{Lip}}(T, T') = \log\left(\frac{\mathrm{covol}(T)}{\mathrm{covol}(T')}\,\mathrm{Lip}(T, T')\right)$$

C'est une pseudométrique asymétrique sur $\mathcal{D}$ (voir la figure 5 pour des exemples) : l'inégalité triangulaire est vraie, mais il existe $T \neq T'$ tel que $d_{\mathrm{Lip}}(T, T') = 0$, et il existe T, T' tel que $d_{\mathrm{Lip}}(T, T') \neq d_{\mathrm{Lip}}(T', T)$.

Lorsqu'un automorphisme $\phi \in \mathrm{Out}(G)$ admet un représentant train track irréductible $f : T \to T\cdot\phi$, celui-ci vérifie $d_{\mathrm{Lip}}(T, T\cdot\phi) = \log(\mathrm{Lip}(T, T\cdot\phi))$. Mieux, on a $d_{\mathrm{Lip}}(T, T\cdot\phi^n) = n\log(\mathrm{Lip}(T, T\cdot\phi))$ et les points $(T\cdot\phi^n)_{n\in\mathbb{Z}}$ sont alignés le long d'une géodésique bi-infinie de $\mathcal{D}$. Remarquons que comme d_{Lip} est asymétrique, un chemin peut être une géodésique lorsqu'on la parcourt dans un sens mais pas dans l'autre sens.

Ainsi un automorphisme irréductible qui admet un représentant train track agit de façon loxodromique sur l'espace de déformation $\mathcal{D}$, et il admet un axe de translation, c'est-à-dire une géodésique bi-infinie invariante de $\mathcal{D}$ sur laquelle ϕ agit par translation.

Dans l'outre-espace, les automorphismes complètement irréductibles, qui admettent

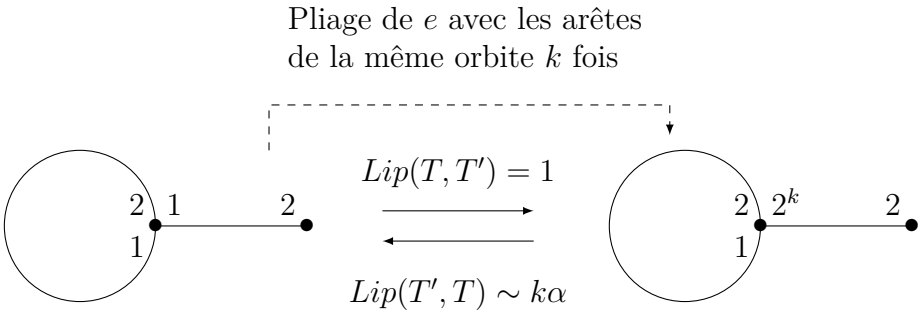

FIGURE 5 – La métrique de Lipschitz est une pseudo-métrique

toujours un train track, ont également un axe de translation. Algom-Kfir a montré dans [Alg11] que l'on peut définir une projection sur un tel axe et que cette projection est *fortement contractante* : soit $\pi_{\mathcal{L}} : \mathrm{CV}_N \to \mathcal{L}$ la projection sur l'axe $\mathcal{L}$ qui à un point $X \in \mathrm{CV}_N$ associe le point $\pi_{\mathcal{L}}(X) \in \mathcal{L}$ qui minimise la distance $d_{\mathrm{Lip}}(X, \pi_{\mathcal{L}}(X))$. Il existe une borne $D > 0$ telle que si B est une boule sortante qui n'intersecte pas $\mathcal{L}$ alors $\operatorname{diam} \pi_{\mathcal{L}}(B) < D$.

Cette propriété permet entre autres de définir la projection d'un axe sur un autre et de montrer une version de l'inégalité de Behrstock pour cette projection. Cela permet d'appliquer la construction de Bestvina-Bromberg-Fujiwara ([BBF15]) pour construire une action de $\mathrm{Out}(F_N)$ sur un *quasi-arbre* qui est, à plus forte raison, un espace hyperbolique au sens de Gromov. Cette application constitue une motivation forte pour trouver un analogue pour les groupes GBS, d'autant plus qu'on ne connaît pas d'action intéressante des groupes d'automorphismes de groupes GBS sur des espaces hyperboliques.

Nous montrons la propriété de forte contraction pour un axe d'automorphisme de groupe de Baumslag-Solitar généralisé :

Théorème O. *Soit G un groupe GBS défini par un graphe Γ avec $b_1(\Gamma) \geq 3$. Soit $\phi \in \mathrm{Out}(G)$ un automorphisme complètement irréductible sans élément pseudo-périodique et $f : T \to T$ un représentant train track pour ϕ, et $f' : T' \to T'$ un représentant train track pour ϕ^{-1}. L'automorphisme ϕ admet un axe $\mathcal{L}_f \subset \mathcal{D}$. Alors on peut définir une projection au point le plus proche $\pi_f : \mathcal{D} \to \mathcal{L}_f$ tel qu'il existe une constante $D > 0$ telle que toute boule sortante $B_{\to}(x, r)$ qui n'intersecte pas $\mathcal{L}_f$ vérifie*

$$\operatorname{diam} \pi_f(B_{\to}(x, r)) \leq D$$

L'automorphisme ϕ^{-1} admet également un axe, qui est généralement différent, mais

les deux axes sont à distance bornée l'un de l'autre.

En revanche, la question d'étendre la propriété aux projections d'un axe sur un autre est encore ouverte. La principale raison à cela est la différence entre la métrique de Lipschitz dans l'outre-espace et dans $\mathcal{D}$. La preuve d'Algom-Kfir utilise une propriété de quasi-symétrie de la métrique de Lipschitz sur la partie épaisse de l'outre-espace. Il n'y a pas de propriété analogue dans $\mathcal{D}$, notamment lorsque l'espace de déformation est de dimension infinie. Il est même parfois possible de trouver T, T' tels que $d_{\mathrm{Lip}}(T, T') = 0$ et $d_{\mathrm{Lip}}(T', T)$ soit arbitrairement grand.

Soit $X \in \mathcal{D}$. On veut définir la projection de Y sur $\mathcal{L}_f$, qui est une géodésique qui relie les points $\dots T \cdot \phi^{-1}, T, T \cdot \phi, T \cdot \phi^2 \dots$

Définition P. Un *candidat* dans X est un élément $g \in G$ tel que l'application $\pi : \mathrm{Axe}_X(g)/\langle g \rangle \to X/G$ prend l'une des formes suivantes, illustrées par la figure 6 :

— une boucle

— un huit : deux cercles plongés qui s'intersectent en exactement un point

— un haltère : deux cercles plongés disjoints reliés par un segment

— un haltère simplement dégénéré : l'un des deux cercles de l'haltère est réduit à un point

— un haltère doublement dégénéré : les deux cercles sont réduits à des points.

Un candidat traverse chaque arête au plus deux fois.

Lemme Q. *Soit X, Y dans $\mathcal{D}$. Il existe un candidat g_X de X tel que*

$$\mathrm{Lip}(X, Y) = \frac{\|g_X\|_Y}{\|g_X\|_X}$$

L'approche pour calculer $d_{\mathrm{Lip}}(X, \mathcal{L}_f)$ consiste donc à évaluer $\frac{\|g_X\|_{T \cdot \phi^n}}{\|g_X\|_X}$ en faisant varier n pour tous les candidats, ces derniers existant en nombre fini.

Étant donné un élément simple $g \in G$, on peut considérer son axe dans chacun des $T \cdot \phi^n$ pour $n \in \mathbb{Z}$ (et chacun des $T' \cdot \phi^n$ également). L'axe de g n'est pas forcément légal et se découpe en segments légaux maximaux. On définit $\mathrm{Leg}_f(g)$ comme la proportion de l'axe de g constituée de segments légaux de longueur supérieure à une certaine constante κ_f. Un fait essentiel est qu'à partir d'un certain rang pour n croissant, $\mathrm{Leg}_f(g)$ dépasse un certain seuil et on a alors $\|\phi^n(g)\|_T \to \infty$ avec une croissance exponentielle et contrôlée

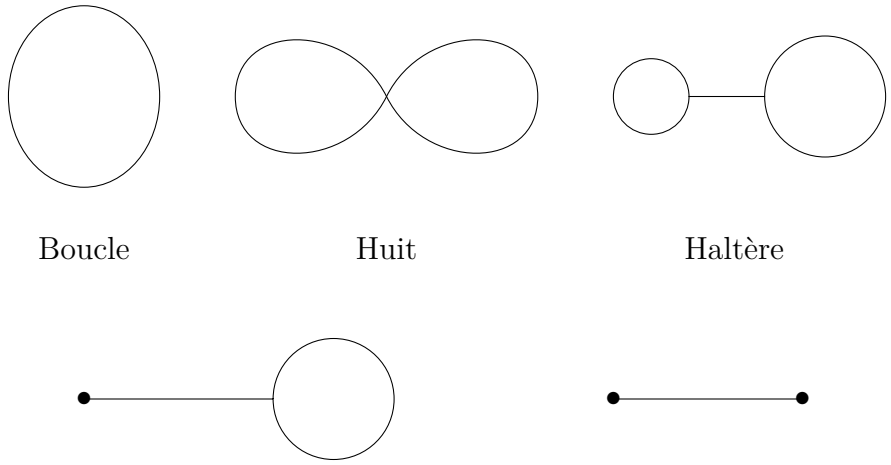

FIGURE 6 – Les cinq formes possibles de candidats

(Lemme 3.5.1). L'idée est que les longs segments légaux prennent de plus en plus de place dans l'axe de g.

On peut définir de même $\mathrm{Leg}_{f'}(g)$. On a le résultat suivant lorsque ϕ n'a pas d'élément pseudo-périodique :

Lemme R. *Il existe une constante $\epsilon_0 > 0$ et $N \in \mathbb{N}$ tels que pour tout élément loxodromique $g \in G$ on ait*

$$\mathrm{Leg}_f(\phi^N(g)) > \epsilon_0 \text{ ou } \mathrm{Leg}_{f'}(\phi^{-N}(g)) > \epsilon_0$$

Ce lemme, prouvé dans [BFH97] pour les groupes libres, permet de montrer que $\{n \in \mathbb{Z} | \|g\|_{T\cdot\phi^n} \text{ minimal}\}$ est de diamètre borné indépendamment de g.

On peut ainsi définir un sous-segment de l'axe de diamètre borné indépendamment de X, que nous appelons $\pi_f(X)$.

Cette projection admet des propriétés de courbure négative, telles que :

Lemme S. *Il existe $s, c > 0$ telles que pour tout $Y \in \mathcal{D}$, pour tout $S \in \mathcal{L}_f$, si $d(\pi_f(Y), S) > s$ alors $d(Y, S) \geq d(Y, \pi_f(Y)) + d(\pi_f(Y), S) - c$.*

Un ingrédient important de la preuve est donné par la proposition suivante. On dit qu'une paire d'éléments $\alpha, \beta \in G$ est *simple* s'il existe un système de facteurs cycliques qui les contient tous les deux. Étant donné $L > 0$, on dit que α *contient un L-morceau*

de la lamination stable $\Lambda^+(\phi)$ si l'axe de α contient un segment de feuille de $\Lambda^+(\phi)$ de longueur au moins L. On a :

Proposition T. *Il existe $L > 0$ tel que pour tout $S \in \mathcal{L}_f$ on ait :*

- *Si $\beta \in G$ contient à la fois un L-morceau de $\Lambda^+(\phi)$ et un L-morceau de $\Lambda^-(\phi)$ alors β n'est pas contenu dans un facteur cyclique*
- *Si $\alpha \in G$ contient un L-morceau de $\Lambda^+(\phi)$ et $\beta \in G$ contient un L-morceau de $\Lambda^-(\phi)$ alors la paire $\{\alpha, \beta\}$ n'est pas simple.*

La preuve de ce résultat utilise les graphes de Whitehead des laminations stable et instable pour donner des informations sur ceux de α et β. Le théorème 1.3.3 permet alors de conclure.

Chapitre 1

Whitehead algorithm for automorphisms of generalized Baumslag-Solitar groups

Baumslag-Solitar groups, defined by $\mathrm{BS}(p,q) = \langle a, t | ta^p t^{-1} = a^q \rangle$, are a family of one-relator groups which were introduced to give examples of non-hopfian groups in [BS62]. Generalized Baumslag-Solitar (GBS) groups are a wider family composed of all fundamental groups of finite graphs of infinite cyclic groups. The isomorphism problem between GBS groups (i.e. determining if two graphs of groups define the same GBS group) is not known apart from a few special cases. Clay and Forester showed in [CF08a] that it is solvable when first Betti number is at most one. Levitt ([Lev07]) defined a particular class of GBS groups with nice outer automorphism group for which the isomorphism problem is also solvable.

In fact, the automorphism group of a GBS group G can be very complicated in general and depends a lot on G. For example, $\mathrm{Out}(\mathrm{BS}(2,3))$ is finite but $\mathrm{Out}(\mathrm{BS}(p,pn))$ for $p, n \in \mathbb{Z}$ is not finitely generated. This was shown by Collins and Levin in [CL83] with algebraic methods and later by Clay in [Cla09] using geometric methods.

By Bass-Serre theory GBS groups have a natural action on trees which is a motivation for using geometric methods. The study of automorphisms of free groups and outer space give some inspiration. Given a GBS group G, there is an analogue for outer space, which is called *deformation space* ([For06]). It can be defined as the set of all minimal actions of G on trees with infinite cyclic stabilizers, up to G-equivariant isometry and $\mathrm{Out}(G)$ acts on it by precomposition of the action.

An automorphism of the free group F_N is called *fully irreducible* if it does not have any periodic conjugacy class of free factor of F_N. Such an automorphism ϕ admits a *train track* representative: there exists a free minimal action of F_N on a simplicial metric tree, which defines a translation axis for the action of ϕ on outer space. Train tracks representatives

are a powerful tool to study automorphisms of F_N.

In some sense one can define fully irreducible automorphisms of a GBS group G: they are the automorphisms whose powers do not preserve any conjugacy class of *special factors*, the analogue of free factors in this context. Margot Bouette showed in [Bou16] that all automorphisms of $\mathrm{BS}(p, pn)$ are reducible, which is surprising: one would have suspected that a generic automorphism would not preserve any special factor. However, she proved that there is a different deformation space invariant under $\mathrm{Out}(G)$ on which irreducible automorphisms exist, always admit train tracks and act on the modified deformation space with positive translation length. In more general GBS groups it is not known if train tracks always exist for fully irreducible automorphisms. Under some restrictions on the GBS group though (ex: no nontrivial integer modulus, [For06]), we can assume that the dimension of the deformation space is finite. Then [Mei15, Theorem 50] applies and proves the existence of train tracks.

This sort of problem is a motivation to understand *special factors*. Here is a way to define them. Identifying the free group F_N with the fundamental group of some graph Γ, the fundamental group of any subgraph of Γ is a free factor of F_N. Moreover, allowing to vary the graph and identification, any free factor of F_N can be obtained this way. Similarly, a special factor of a GBS group G is the fundamental group of some subgraph of a graph of groups Γ where Γ is a graph of cyclic groups with fundamental group G. To avoid degenerate cases, we do not consider cyclic subgroups as special factors. A noticeable difference is the fact that many different graphs of groups may appear and give distinct special factors, while all free factors of free groups may be seen in roses. Therefore, while free factors of same rank in a free group F_N are isomorphic and in the same orbit under $\mathrm{Aut}(F_N)$, special factors of a GBS group may be a lot more diverse and there may be infinitely many orbits of special factors under $\mathrm{Aut}(G)$. This is the case for $G \simeq \mathrm{BS}(2, 4)$.

The Whitehead algorithm in the free group solves the following problem: given $g, h \in F_N$, is there $\phi \in \mathrm{Aut}(F_N)$ such that $\phi(g) = h$? A weaker form of this algorithm can decide, given $g \in F_N$, if g is simple, i.e. whether there exists a proper free factor containing g. The first problem seems difficult for GBS groups since the automorphism group is a lot more complicated and may not be finitely generated. Here we solve the analogue of the second problem in a GBS group G by giving an algorithm deciding if an element $g \in G$ is contained in a special factor. We also prove that there exists a unique minimal special factor containing g and give a further algorithm which finds this minimal factor.

Before stating the results, let us give some useful background. Let G be a GBS group.

It admits an action on a locally finite tree T with cyclic edge and vertex stabilizers. Unless G is isomorphic to $\mathbb{Z}, \mathbb{Z}^2$ or the fundamental group of a Klein bottle, all trees with cyclic stabilizers have the same elliptic subgroups. Let $\mathcal{D}$ be the set of G-trees with cyclic edge and vertex stabilizers. We call it the *cyclic deformation space*.

In the main part of the paper we will consider *restricted deformation spaces*, where we impose an extra condition on the edge groups of the trees. In this introduction we present the result for $\mathcal{D}$ the cyclic deformation space for simplicity but we actually prove a slightly stronger version.

A *loxodromic* element in a tree T in $\mathcal{D}$ is an element acting on a tree with no fixed point. This property does not depend on the choice of $T \in \mathcal{D}$, so we may refer to loxodromic elements of the group G. A loxodromic element of G is *simple* if there exists a proper special factor containing it. In order to understand whether an element is simple, we use *Whitehead graphs*. The Whitehead graph $\mathrm{Wh}_T(g, v)$, where v is a vertex, has vertices corresponding to the edges of T with origin v. Two vertices are linked by an edge whenever a translate of the axis of g in T takes the turn between the two edges corresponding to the vertices. We say that the graph has an *admissible cut* when it is disconnected or has a cut vertex.

Whitehead's lemma was originally published in [Whi36] for free groups. We have the following result, adapted from a version of Whitehead's lemma from [GH19].

Theorem A. *Let $g \in G$ a loxodromic element. If g is simple then for all $T \in \mathcal{D}$ there exists a vertex $v \in T$ such that $\mathrm{Wh}_T(g, v)$ has an admissible cut.*

Let us take an element $g \in G$ and a tree $T \in \mathcal{D}$. We consider T as a marked graph of groups $\Gamma = T/G$. In case $\mathrm{Wh}_T(g, v)$ has an admissible cut, either g is represented in the fundamental group of some proper subgraph of Γ, so it is simple, or we can perform some transformation of Γ in order to obtain a new tree. We can apply the theorem above to the result of the transformation. The point is that this process will eventually stop, so we deduce an algorithm to determine whether g is simple:

Corollary B. *There is an algorithm taking as input*

— *a GBS group G given as a graph of groups*

— *a loxodromic element $g \in G$*

which decides whether g is simple, and if it is, returns a proper special factor containing g.

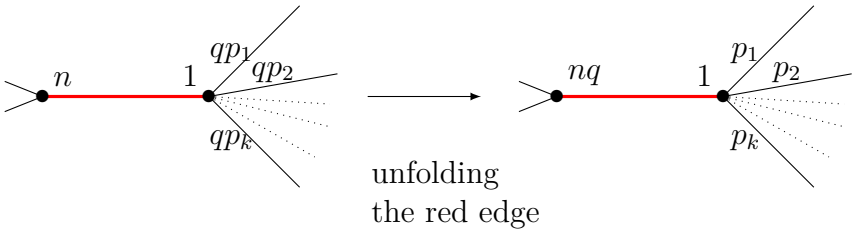

Figure 1.1 – Unfolding of type II (i.e. the unfolded edges belong to the same orbit) as seen in the graph of groups

The point of the algorithm will be to transform the graph until one of its proper subgraphs contains the axis of g. Note that Whitehead's version for free groups (see [LS01, Chapter I, section 4]) has the same goal but uses specific automorphisms to transform a rose. Here we prefer to use the approach given in [GH19] instead. It consists in *unfolding* the tree associated to the graph, that is to say, perform the inverse of folds. While in Whitehead's approach the word length decreases and the graph keeps the same volume, here the word length remains constant while the volume of the graph increases, eventually reaching a point where g avoids some edge in the graph.

The main reason for which we choose this approach is that there is no preferred graph of groups in general, and the graph of group showing the special factor containing g might be very different from the initial graph of groups used to define G. There might be no automorphism between the factor containing g and a subgraph of the initial graph.

We prove a stronger version of Theorem A and Corollary B which applies to a finite collection of elements of G rather than a single element. Actually the version for collections decides whether the elements of a collection belong to a *system of special factors*, which are special factors which are either conjugate or disjoint in some sense.

In the process of showing that the algorithm stops, we need to show that the volume of the graph increases. It does not increase at each step and in fact, if the transformation performed is an *unfolding* of edges in the same orbit (see Figure 1.1), the volume remains constant. The following result prevents the graph from being transformed indefinitely without increasing its volume.

A *sequence of unfoldings of type II* is a finite or infinite sequence $T_n \to \cdots \to T_1 \to T_0$ where the T_i are trees in $\mathcal{D}$ and $T_{i+1} \to T_i$ is a type II fold, that is, a fold of edges in the same orbit, with origin and endpoint in different orbits. Then no such sequence can be infinite:

Lemma C. *Suppose that G is not a solvable Baumslag-Solitar group. Any sequence of unfoldings of type II is finite.*

Whitehead graphs, which are useful for the algorithm, can be computed algorithmically. In [Bee15] methods for algorithmic computations in GBS groups are given. In section 1.3 we explain the algorithm in detail.

As a special factor is itself a GBS group, we may want to iterate the algorithm in the following way. Given $g \in G$ belonging to a special factor $H \in G$, we can apply the algorithm to H in order to find a smaller special factor, again and again. In the case of free groups, the rank of the free factor decreases, so eventually we find a free factor which does not have any smaller free factor containing the element. In the case of GBS groups the rank does not always decrease, but we do have a complexity $\mathcal{C}$ which decreases strictly when passing to a proper special factor. Note that in some GBS groups including $\mathrm{BS}(2,4)$ there exist arbitrary long chains of decreasing special factors (see Figure 1.2). For the definition of $\mathcal{C}$, see section 1.4.

Proposition D. *Let G be a GBS group and H a special factor of G. Then $\mathcal{C}(H) < \mathcal{C}(G)$.*

From this we deduce the existence of a minimal special factor containing g, which is in fact unique:

Theorem E. *The set of special factors relative to $\mathcal{D}$ which contain a given loxodromic element g admits a smallest element for inclusion.*

Using the previous algorithm repeatedly we get the following:

Theorem F. *There exists an algorithm taking as input a GBS group G as a marked graph of groups and a hyperbolic element g, and which outputs the minimal factor containing g as a subgraph of a marked graph of groups for G.*

Just like Theorem A, Theorem F also applies to finite collections of elements of G. In that case it outputs the minimal system of special factors containing the collection.

The paper is structured as follows. In section 1.1 we introduce basic notions about GBS groups and prove that no sequence of unfoldings of type II can be infinite. In section 1.2 we prove Theorem A. Section 1.3 proves Theorem F: it gives of an algorithm to determine whether there exists a special factor containing a particular group element. Most of it consists in proving why all steps needed to check that an element is simple are algorithmic. Since it is more technical, the proof of the algorithm may be skipped at first reading. Finally we prove theorem E in section 1.4.

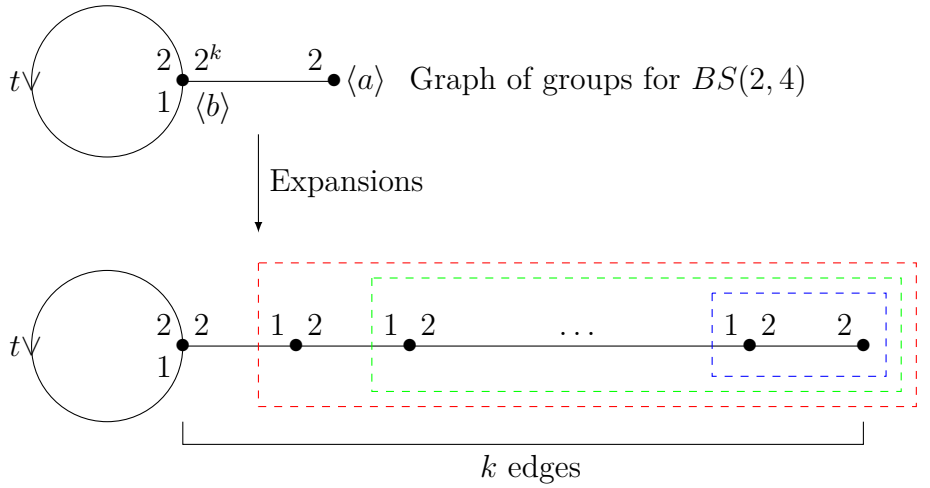

Figure 1.2 – A construction of arbitrary long special factor sequences in BS$(2,4)$. The upper graph represents BS$(2,4)$ with the presentation $\langle a, b, t | tbt^{-1} = b^2, a^2 = b^{2^k} \rangle$. Expansions from this graph lead to the graph below, which has many edges. The rectangles show some special factors which can be read in the graph below. Here we can construct a sequence of k nested special factors.

1.1 Definitions

1.1.1 Trees, elliptic groups, deformation spaces

We refer to [Ser77] for basic notions on graphs of groups. A *graph* is a set of vertices, along with a set of oriented edges. An edge comes with applications o and t which associate its initial and terminal vertex to an edge, respectively. There is a fixed-point-free involution $e \mapsto \bar{e}$ with $t(\bar{e}) = o(e)$.

A *graph of groups* is a graph Γ such that every vertex v is labelled with a group G_v, every edge e is labelled with a group $G_e = G_{\bar{e}}$ and for every oriented edge e of Γ, there is a given monomorphism $\phi_e : G_e \to G_{t(e)}$.

Definition 1.1.1. A *generalized Baumslag-Solitar group* (*GBS* group) is the fundamental group of some finite graph of groups of which each edge or vertex group is infinite cyclic.

A GBS group is *non-elementary* if it is not isomorphic to one of the following: $\mathbb{Z}$, $\mathbb{Z}^2$ or the fundamental group of the Klein bottle $\langle a, t | tat^{-1} = a^{-1} \rangle = \langle a, b | a^2 = b^2 \rangle$.

Remarks 1.1.2. — Generally the graph of (cyclic) groups for a GBS group is not unique, in fact there may be infinitely many such graphs.

— From a graph of groups one can deduce a presentation for the group. Since vertex

groups are cyclic, GBS groups are finitely presented. See [Lev07] for some detail on GBS groups.

Definition 1.1.3. A *labelled graph* is a graph of which each oriented edge e carries a label $\lambda(e) \in \mathbb{Z} \setminus \{0\}$ at its origin.

In graphs of cyclic groups, choosing a generator a_v for each vertex group and a_e for each edge group gives an identification with $\mathbb{Z}$ and each inclusion $G_e \to G_v$ is given by a multiplication by a nonzero integer. Therefore any such graph of groups can be described by a labelled graph.

Definition 1.1.4. An edge in a graph is a *loop* if its endpoints are equal.

Let G be a non-elementary GBS group.

Definition 1.1.5. A *G-tree* is a simplicial tree endowed with a minimal action of G by simplicial isomorphisms without inversion of edges.

We endow all trees with the combinatorial metric, that is, all edges have length 1.

Remark 1.1.6. Let T be a G-tree and let g be a loxodromic element in T, that is, an element acting on T with no fixed point. The *translation length* of g is

$$\|g\|_T := \min_{x \in T} d(x, gx)$$

and it is equal to of edges in a fundamental domain of the axis of g in T.

Let T be a G-tree. Let e be an edge of T. We denote the stabilizer of e by G_e (and this subgroup fixes both endpoints of e since the action is without inversion). Similarly we denote by G_v the stabilizer of a vertex v.

Definition 1.1.7. Let T be a G-tree and H a subgroup of G. The subgroup H is *elliptic* (resp. *bi-elliptic*) in T if it fixes a vertex (resp. an edge).

Definition 1.1.8. Two trees T, T' are in the same *deformation space* if they share the same elliptic subgroups. If a subgroup $H < G$ fixes a point in T, it must then fix a point in T', and conversely.

Remark 1.1.9. Equivalently, T and T' belong to the same deformation space if there exists G-equivariant applications $T \to T'$ and $T' \to T$ (see [GL07] for detail).

Apart from $\mathbb{Z}, \mathbb{Z}^2$ and the fundamental group of the Klein bottle, all actions on trees dual to descriptions of G as graphs of cyclic groups belong to the same deformation space ([For06]).

Definition 1.1.10. For a non-elementary GBS group G, the *cyclic deformation space* $\mathcal{D}$ is the space of all G-trees with infinite cyclic vertex and edge stabilizers.

We define a restricted deformation space of the cyclic deformation space, which is smaller than $\mathcal{D}$, like in [GL07]:

Definition 1.1.11. Let $\mathcal{A}$ be a family of subgroups of G, stable by conjugation and by taking subgroups. The *restricted deformation space* $\mathcal{D}^{\mathcal{A}}$ is the set of G-trees in $\mathcal{D}$ whose bi-elliptic subgroups all belong to $\mathcal{A}$.

We refer to subgroups which belong to $\mathcal{A}$ as *allowed edge groups* in $\mathcal{D}^{\mathcal{A}}$.

Remark 1.1.12. For a subgroup of G, being elliptic is a property which depends only on the deformation space of the tree studied. However in general this does not hold for bi-elliptic subgroups: two trees in the same deformation space may have different bi-elliptic subgroups. This applies to $\mathcal{D}$ and $\mathcal{D}^{\mathcal{A}}$.

Definition 1.1.13. Let be T a G-tree. Let F be a G-invariant subforest of T. We define the equivalence relation $\sim_F$ on T as the smallest G-invariant equivalence relation such that $x \sim_F y$ whenever there exists some connected component of F containing both x and y.

The quotient $T \to T/\sim_F$ is a tree, called the *collapse* of F in T.

The G-equivariant application $T \to T/\sim_F$ is called the *collapse map*.

When F is the G-equivariant subforest spanned by an edge e, we speak of the *collapse of* e instead of the collapse of F, and denote it by $T \to T/\sim_e$

Definition 1.1.14. Let T be a G-tree. An edge of T is called *collapsible* if $T/\sim_e$ is a tree in the same deformation space as T. An edge e with endpoints u and v is collapsible if u and v are in different orbits and either $G_e = G_u$ or $G_e = G_v$.

Remark 1.1.15. If there is a G-equivariant application $f : T \to S$ where $T, S \in \mathcal{D}$ such that the image of e is a single point in S, then e is collapsible. Indeed f factorizes through $T/\sim_e$. Thus there exists a G-equivariant application $T/\sim_e \to T$ (and vice-versa) so $T/\sim_e$ is in $\mathcal{D}$.

Definition 1.1.16. A G-tree T is *reduced* if no edge in T is collapsible.

Remarks 1.1.17. 1. This notion can be expressed of in terms graphs of groups, which are more convenient for a computational use. Denote by Γ the graph of groups associated to the quotient T/G, then T is reduced if whenever an edge morphism in Γ into a vertex group is surjective then the edge is a loop.

2. Reduced trees in a deformation space share the same bi-elliptic subgroups. Any bi-elliptic subgroup in some reduced tree $T \in \mathcal{D}$ is also bi-elliptic in any tree in $\mathcal{D}$, reduced or not.

Definition 1.1.18. We call $\mathcal{A}_{\min}$ the family of subgroups which are bi-elliptic in reduced trees. The *reduced deformation space* $\mathcal{D}_{\text{red}}$ is the restricted deformation space $\mathcal{D}^{\mathcal{A}_{\min}}$.

Definition 1.1.19. We call an elliptic subgroup in some tree $T \in \mathcal{D}$ *big* with respect to $\mathcal{A}$ if it is not in $\mathcal{A}$.

Remark 1.1.20. In order to check whether an elliptic subgroup is big for $\mathcal{A}_{\min}$, it suffices to check that it fixes no edge in some arbitrary tree in $\mathcal{D}$.

In a graph of cyclic groups viewed as a labelled graph, vertex groups such that no label at the vertex is ± 1 are maximal big groups. If the graph of groups is reduced, then maximal big groups are exactly vertex stabilizers with all labels different from ± 1.

The number of conjugacy classes of maximal big groups is finite (bounded by the number of vertex orbits of some tree in $\mathcal{D}$) and depends only on G and $\mathcal{D}$ (see [GL07]).

1.1.2 Folds, expansions

In this section we define folds and expansions, and give a construction for a certain type of expansion.

Definition 1.1.21. A G-tree S is a *refinement* of another G-tree T if there exists a collapse $\pi : S \to T$, i.e. T is equivariantly isomorphic do $S/\sim_F$ for some G-invariant forest $F \subset S$.

We say that S is an *expansion* if additionnally S and T belong to the same deformation space.

Lemma 1.1.22 (Construction of an expansion $T^{H,S}$)**.** *Let T be a G-tree and v a vertex of T. Let $\mathcal{E}_v$ be the set of edges with origin v. Let H be a subgroup of G_v and S a non-empty proper subset of $\mathcal{E}_v$ which satisfy:*

— $HS = S$

— $\forall g \in G_v \setminus H,\ gS \cap S = \varnothing$.

Consider the partition $\mathcal{E}_v = \bigsqcup_{g \in G_v/H} gS \sqcup E'$ *where* $E' = \mathcal{E}_v \setminus G_v \cdot S$.

We construct $T^{H,S}$ *from* T *as follows. First we replace* v *by the star on* G_v/H *and for each* $g \in G_v$ *and each edge* $e \in S$*, we attach* $g \cdot e$ *to* gH*. Then we attach* E' *on the centre of the star. Finally we extend this by equivariance to all translates of* v*.*

The obtained tree $T^{H,S}$ *is minimal, is a refinement of* T *and belongs to the same deformation space (without restriction on edge groups). Moreover, if* $H \in \mathcal{A}$ *and* $T \in \mathcal{D}^{\mathcal{A}}$ *then* $T^{H,S} \in \mathcal{D}^{\mathcal{A}}$*.*

Proof. Let us prove that $T^{H,S}$ is minimal. The tree $T^{H,S}$ has no valence 1 vertex: since $S \subsetneq \mathcal{E}_v$ and $S \neq \varnothing$ then either G_v/H has two or more elements, or E' is non-empty. In both cases no vertex has valence 1.

If $T^{H,S}$ were not minimal, it would have a valence 1 vertex: suppose $T^{H,S}$ is not minimal and call $T^{H,S}_{\min}$ the minimal invariant subtree. By cocompacity of the action there exists $v \in T$ such that the distance $d(v, T_{\min})$ is maximal and positive. Then v has valence 1, which gives the contradiction needed.

Note that collapsing the stars involved in the construction of $T^{H,S}$ allows to recover the original tree T, so $T^{H,S}$ is a refinement of T. This implies that all elliptic subgroups of $T^{H,S}$ are elliptic in T. Conversely, let G_0 be a subgroup fixing a vertex w in T. If $w \notin G \cdot v$ then its pre-image by the collapse map is a single point so it must be fixed by G_0. If $w = v$, its pre-image is the closure of the star. The star is invariant by G_v so its center is a fixed point for G_v which contains G_0. This proves that T and $T^{H,S}$ have the same elliptic subgroups.

Thus $T^{H,S}$ is in the same deformation space as T. Note that the second condition implies that the stabilizer of any edge of the star is conjugate to H, hence the last statement. □

To define folds, we rely on [BF91]. Folds are a classical operation on trees and were first defined by Stallings ([Sta91]). The idea behind a fold is to identify two edges of a G-tree in an equivariant way to create a new tree.

Definition 1.1.23 (Fold)**.** Let T be a G-tree. Let V be a vertex of T and e_1, e_2 two edges with origin v such that $e_2 \notin G \cdot \bar{e}_1$. We define the fold of e_1 together with e_2 as follows.

We define the equivalence relation $\sim$ on T as the smallest G-invariant equivalence relation satisfying $e_1 \sim e_2$ and $t(e_1) \sim t(e_2)$.

The result of the fold of e_1 with e_2 is $T/\sim$ and it is a tree.

Up to subdividing some edges, any fold boils down to a sequence of folds of the three types illustrated by Figure 1.3 (see [BF91] for more details on fold types), which presents folds as seen in the quotient graph, depending on whether the edges an their terminal vertices are in the same orbit or not.

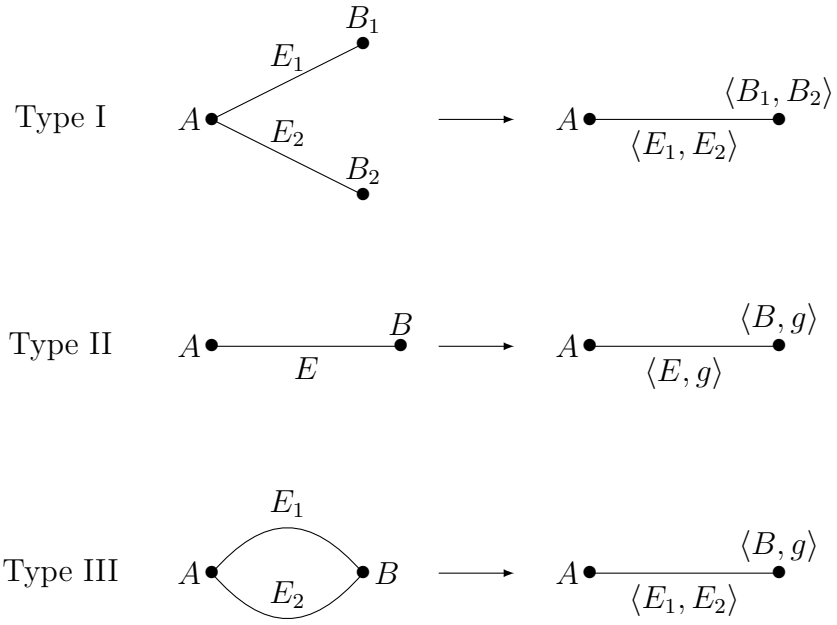

Figure 1.3 – The three basic types of folds

Denote by $A = G_v$, $B_i = G_{t(e_i)}$, $E_i = G_{e_i}$. The result of a fold is in the same deformation space as the original tree if and only if one of the following conditions [1] is true (see figure 1.4):

— the fold is of type I with $B_1 = E_1$ and $B_2 = E_2$, which we call *type A* (not referring to subtypes of [BF91])

— the fold is of type I and $B_2 \subset B_1$ or $B_1 \subset B_2$ which we call *type B*

— the fold is of type II and $E_1 = B_1$ which we call *type C*.

Remark 1.1.24. The pre-image of any edge by a fold which does not change the deformation space is star-shaped, i.e. it consists in a collection of edges which share a common endpoint. In fact

1. If we work with restricted deformations spaces $\mathcal{D}^{\mathcal{A}}$, we must ensure when folding that the new edge groups belong to $\mathcal{A}$. However in this article we will work backwards: starting with the folded tree, we will unfold it, and we do not need to worry since the new groups will be smaller.

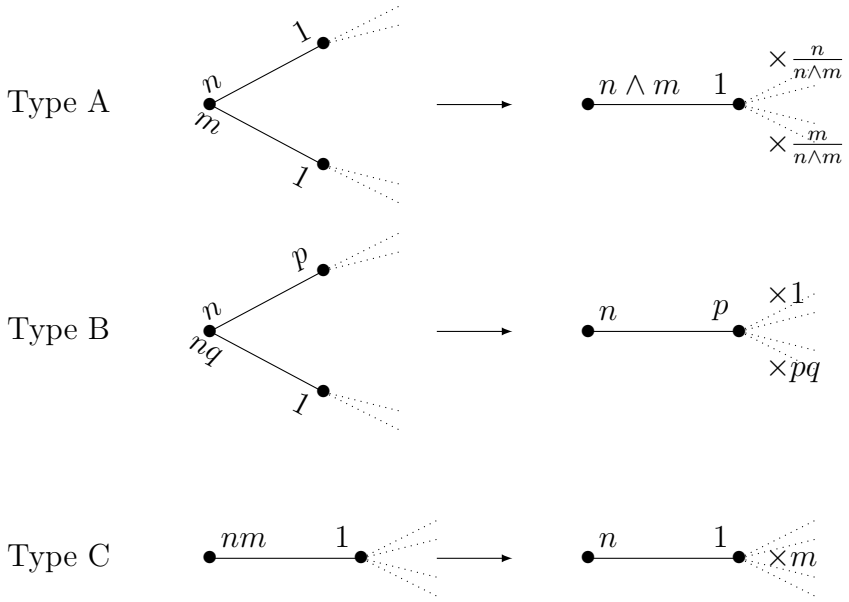

Figure 1.4 – The three types of folds which preserve the deformation space, as seen in labelled graphs.

— $\pi^{-1}\pi(\epsilon) = \epsilon$ if $\epsilon \notin G \cdot e_1, e_2$

— $\pi^{-1}\pi(e_1) = G_{\pi(e_1)} \cdot \{e_1 \cup e_2\}$ and $G_{\pi(e_1)}$ fixes $o(e_1) = o(e_2)$ in case A, B and C.

An *unfolding of type C* is the inverse of a fold of type C, i.e. given a tree T_0, it consists in finding a tree T_1 such that $T_1 \to T_0$ is a fold of type C. We will need the following result about sequences of unfoldings. A *sequence of unfoldings* of type C is a sequence $\cdots \to T_n \to \cdots \to T_1 \to T_0$ such that every $T_{i+1} \to T_i$ is a fold of type C.

Lemma 1.1.25. *Suppose that G is not a solvable Baumslag-Solitar group. Let T be a GBS tree for G. Any sequence of unfoldings of type C (see Figure 1.3) of T is finite.*

Proof. We will show that it is impossible to construct an infinite sequence of unfoldings of type C. We consider the quotient $\Gamma = T/G$. Unfoldings of type C do not change the edges of Γ: only labels vary. Therefore in the rest of the proof edges will keep the same name after unfolding. Consider the product $\prod_{e \in T/G} |\lambda(e)|$ of all labels in Γ (see Figure 1.5). It is a positive integer.

An unfolding of type C can be performed on an edge $e \in \Gamma$ if and only if $\lambda(\bar{e}) = \pm 1$ and there exists an integer q such that $|q| > 1$ and q divides all other labels at $t(e)$.

During the unfolding, the labels at only two vertices may change : $\lambda(e'), e' \in \mathcal{E}_{t(e)} \setminus \{e\}$ are divided by q whereas $\lambda(e)$ is multiplied by q. Let k be the valence of the vertex $t(e)$,

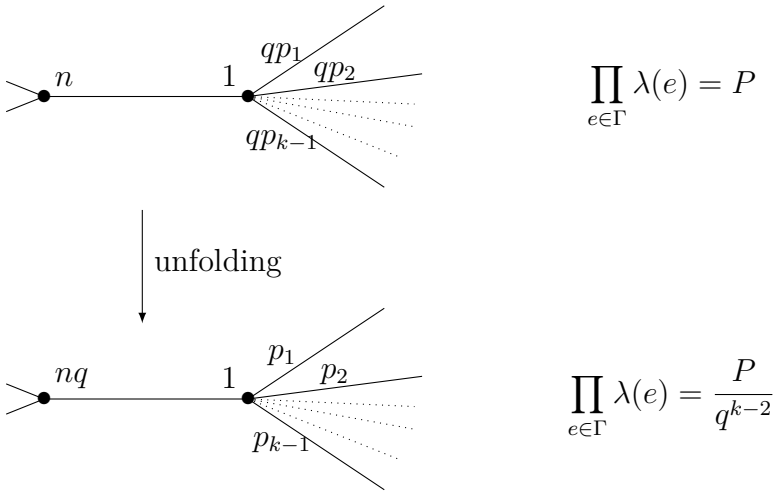

Figure 1.5 – The product decreases when doing unfoldings of type C.

where $k \geq 2$ by minimality. The product $\prod_{e\in T/G} |\lambda(e)|$ is multiplied by q^{2-k}. If the valence of $t(e)$ in Γ is at least 3 (i.e. $k \geq 3$) then the product of labels decreases. An unfolding at a valence 2 vertex does not change the product. To conclude we need to show that one cannot produce an infinite sequence of unfoldings of type C on edges whose terminal vertex has valence 2.

In the rest of this proof we will denote this sort of unfolding by "unfolding from a valence 2 vertex". We will say that the unfolding of an edge e is an unfolding "from $t(e)$". Since no new edge is created in Γ, edges will keep the same name after unfolding.

The *topological edges* of the graph Γ are connected components of the graph without its vertices with valence greater or equal to 3. Since the number of topological edges of the graph is finite any infinite sequence of unfoldings from valence 2 vertices would have infinitely many unfoldings in at least one topological edge.

Note that an unfolding from a valence 2 vertex does not change any label outside of its topological edge. Consequently valence 2 unfoldings in different topological edges commute.

In most cases a topological edge is a segment c. It may be a circle, only if Γ is a circle. First let us study the case of a segment (Figure 1.6). It is composed of a certain number of edges with labels, and is bounded by vertices of valence 1 or at least 3. We will show that there is a bound (depending on Γ) on the length of any sequence of unfoldings from inner vertices in the same topological edge. As unfoldings in different topological edges commute, this will show that no infinite sequence of unfoldings from valence 2 vertices

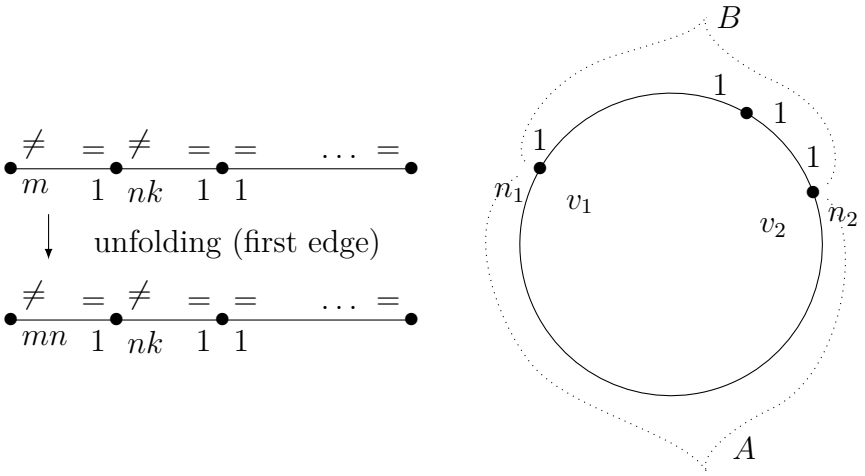

Figure 1.6 – Topological edges

exist.

Choose an orientation for c and view it as a concatenation of oriented edges $e_1 \dots e_l$. We define the following complexity for a segment c of length l:

$$K(c) := \prod_{i=1}^{n} |\lambda(e_i)|^i \times \prod_{i=1}^{n} |\lambda(\bar{e}_i)|^{n+1-i}$$

which gives more weight to initial labels of edges at the end of the chain and to terminal labels of edges at the beginning of the chain. Note that it is well-defined since it does not depend on the choice of orientation for c.

The idea is that unfoldings move factors of labels in the direction where they will weight less. Let us show that the complexity decreases during any unfolding.

Let $i \in \{1, \dots, l-1\}$. Suppose we do an unfolding of e_i (a valence 2 unfolding, so we exclude the case $i = l$). We obtain a new segment c' with different labels. Two labels change: $\lambda(e_{i+1})$ is divided by a factor q and $\lambda(e_i)$ is multiplied by this same factor q, where $|q| \neq 1$. Therefore $K(c') = K(c) \times |q|^i / |q|^{i+1} = K(c)/|q|$ so $K(c') < K(c)$.

Since the formula for the complexity does not depend on the orientation of the segment we get the same result for an unfolding of a $\bar{e}_i$.

As K is a positive integer it cannot decrease indefinitely. This proves that the sequence of unfoldings must stop when the topological edge is a segment.

When the graph is a circle, it has a single topological edge. If there exists a vertex group which is strictly bigger than all edge groups at this vertex, then no unfolding can

occur from this vertex and labels at this vertex can only increase. The graph minus this vertex is a segment to which we can apply the argument above: no unfolding sequence can be infinite.

If not then every vertex has at least one label which is ± 1. Choose an orientation of the circle and orient the edges accordingly. If all initial labels (for this orientation) are ± 1 then the group is a solvable Baumslag-Solitar group. The same deduction can be made with the reverse orientation. If this happens for neither orientation then we can find two vertices with a label different from ± 1 pointing in different directions. Call v_1, v_2 these two vertices, which split the circle into two segments. Call B the segment with the ± 1 labels at its endpoints and A the segment with greater labels. Up to taking a subsegment of B we may suppose that the labels borne by the edges of B are all ± 1.

Unfoldings in A cannot be done indefinitely, because no unfolding from v_1 or v_2 occur and we are again in the case of a segment. This shows that any long enough sequence of unfoldings must involve an edge in B.

The only unfoldings which may involve edges in B are unfoldings from v_1 or v_2. If we perform such an unfolding we increase one of the labels inside B. Then we can define a new partition where B strictly decreases, and iterate until there is a vertex with both labels different from ± 1. □

1.2 Special factors, and algorithm for simplicity of group elements

1.2.1 Special factors for a deformation space

In this section, we introduce special factors with respect to some deformation space $\mathcal{D}^{\mathcal{A}}$. They are an analogue of free factors for free groups.

Remark 1.2.1. We will write $\bar{\mathcal{D}}^{\mathcal{A}}$ to denote the set of all trees obtained by collapsing G-invariant subforests in trees of $\mathcal{D}^{\mathcal{A}}$, including the trivial tree. We include trees of $\mathcal{D}^{\mathcal{A}}$ which correspond to collapsing empty forests.

Definition 1.2.2. A *special factor H with respect to $\mathcal{D}^{\mathcal{A}}$* is a subgroup of G which is the stabilizer of a point in a tree T in $\bar{\mathcal{D}}^{\mathcal{A}}$ and which is not elliptic in $\mathcal{D}^{\mathcal{A}}$.

When the deformation space is obvious, we will write simply *special factor*.

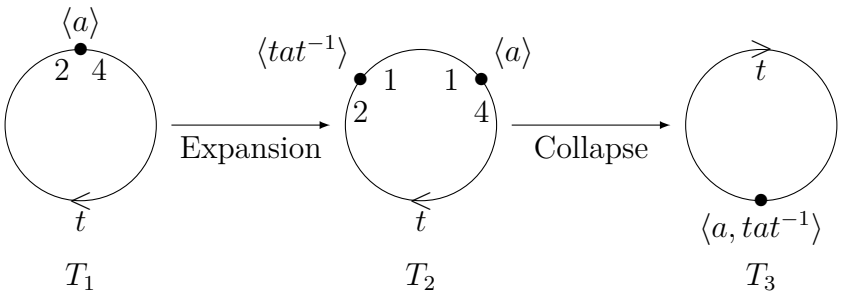

Figure 1.7 – Expansion and collapse on the standard tree T_1 (seen in the quotient) and new special factor with respect to $\mathcal{D}$ obtained after the collapse, visible in $T_3 \in \bar{\mathcal{D}}$. The tree T_2 is in $\mathcal{D}$ but not in $\mathcal{D}^{\mathcal{A}_{\min}}$.

We call H a *proper special factor* when $H \neq G$. Elliptic groups, i.e. vertex stabilizers of trees in $\mathcal{D}^{\mathcal{A}}$, are not considered to be actual special factors.

Examples 1.2.3. This notion depends on the allowed edge groups in $\mathcal{D}^{\mathcal{A}}$. The space $\mathcal{D}^{\mathcal{A}_{\min}}$ of reduced trees has fewer allowed edge groups than $\mathcal{D}$. Consider the standard tree T_1 for $\mathrm{BS}(2,4) := \langle a, t | ta^2t^{-1} = a^4 \rangle$ (see Figure 1.7) and perform an expansion (yielding T_2) and a collapse (yielding T_3) as described by the figure. We obtain a special factor with respect to $\mathcal{D}$, which is the subgroup $\langle a, tat^{-1} \rangle$. This subgroup cannot be obtained by collapsing a tree in $\mathcal{D}^{\mathcal{A}_{\min}}$, thus it is not a special factor with respect to $\mathcal{D}^{\mathcal{A}_{\min}}$.

For a subgroup $H < G$, denote its conjugacy class by $[H]$.

Definition 1.2.4. A *system of special factors with respect to* $\mathcal{D}^{\mathcal{A}}$ is a finite collection of conjugacy classes of subgroups $\mathcal{H} := \{[H_1], \dots, [H_k]\}$ of G such that there exists $T_{\mathcal{H}} \in \bar{\mathcal{D}}^{\mathcal{A}}$ such that $\mathcal{H}$ is the set of conjugacy classes of vertex stabilizers in $T_{\mathcal{H}}$ which are not elliptic in $\mathcal{D}$.

The system is *proper* if it is not $\{[G]\}$.

Remark 1.2.5. Just like special factors, a system of special factors can be viewed in a graph of groups. Is is given by a collection of disjoint subgraphs $\Gamma_1, \dots, \Gamma_k$ of Γ such that for every $i \in \{1, \dots, k\}$ the subgroup H_i is isomorphic to $\pi_1(\Gamma_i)$.

Definition 1.2.6. Let $\mathcal{H} := \{[H_1], \dots, [H_k]\}$ be a system of special factors. We say that a collection $\mathcal{G}$ of elements of G is $\mathcal{H}$*-peripheral*, which we write $\mathcal{G} \preceq \mathcal{H}$, if for any $g \in \mathcal{G}$ there exists $1 \leq i \leq k$ such that g is contained in a conjugate of H_i.

The collection $\mathcal{G}$ is *simple* if there exists a proper system of special factors $\mathcal{H}$ such that $\mathcal{G} \preceq \mathcal{H}$.

A system of cyclic factors $\mathcal{H}'$ is $\mathcal{H}$-*peripheral* ($\mathcal{H}' \preceq \mathcal{H}$) if for every conjugacy class $[H'] \in \mathcal{H}'$ there exists $[H] \in \mathcal{H}$ such that H' can be conjugated into a subgroup of H.

Remarks 1.2.7.
1. Equivalently $\mathcal{H}' \preceq \mathcal{H}$ if there exists a G-equivariant map $T_{\mathcal{H}'} \to T_{\mathcal{H}}$ where $T_{\mathcal{H}'}, T_{\mathcal{H}}$ are defined as in Definition 1.2.4.
2. The relation $\preceq$ defines an order on the set of systems of special factors. It is obviously reflexive and transitive. Suppose $\mathcal{H}' \preceq \mathcal{H}$ and $\mathcal{H} \preceq \mathcal{H}'$, then we get two maps $T_{\mathcal{H}'} \to T_{\mathcal{H}}$ and $T_{\mathcal{H}} \to T_{\mathcal{H}'}$. This implies that $T_{\mathcal{H}'}$ and $T_{\mathcal{H}}$ have the same elliptic subgroups. In particular they have the same non-cyclic vertex stabilizers, so $\mathcal{H} = \mathcal{H}'$. Thus $\preceq$ is antisymmetric.

1.2.2 Whitehead graph and criterion for simplicity

We fix a collection of cyclic allowed edge groups $\mathcal{A}$ and consider a restricted deformation space $\mathcal{D}^{\mathcal{A}}$.

Definition 1.2.8. Let T be a G-tree. A *turn* in T is an unordered pair of distinct edges with same origin. If e, e' are two such edges, the corresponding turn is denoted by $\{e, e'\}$.

When $e = e'$ we call the pair a *degenerate turn*.

A geodesic γ *crosses* a turn $\{e, e'\}$ if γ contains $e \cup e'$.

Let $\mathcal{G} := \{g_1, \dots, g_k\} \in G$ be a finite collection of loxodromic elements. Let $T \in \mathcal{D}^{\mathcal{A}}$ and v a vertex in T. The set $\mathcal{E}_v$ is the set of edges of T with origin v.

Definition 1.2.9. The *Whitehead graph* $\mathrm{Wh}_T(\mathcal{G}, v)$ is the following graph. The vertex set is $\mathcal{E}_v$. Two vertices e, e' are linked by a non-oriented edge in $\mathrm{Wh}_T(\mathcal{G}, v)$ when there exists $g \in \mathcal{G}$ and some conjugate of g whose axis crosses the turn $\{e, e'\}$.

Remarks 1.2.10.
1. Equivalently, we link e and e' by an edge whenever there is $h \in G$ such that $\{he, he'\}$ is a turn in the axis of some $g \in \mathcal{G}$.
2. The Whitehead graph is a simplicial graph. In particular it does not have any loop.
3. The group G_v has a natural action on $\mathrm{Wh}_T(\mathcal{G}, v)$.

Definition 1.2.11. When $\mathrm{Wh}_T(\mathcal{G}, v)$ is not connected, we call *admissible connected component* any connected component in $\mathrm{Wh}_T(\mathcal{G}, v)$ whose stabilizer is in $\mathcal{A}$. When all edge groups are allowed all connected components are automatically admissible.

Let p be a vertex in $\mathrm{Wh}_T(\mathcal{G}, v)$. Let W_0 be the connected component which contains p. The vertex p is an *admissible cut point* if $W_0 \setminus \{p\}$ is disconnected and if there exists a connected component A of $W_0 \setminus \{p\}$ satisfying $A \cap G_v \cdot p = \varnothing$.

The Whitehead graph $\mathrm{Wh}_T(\mathcal{G}, v)$ has an *admissible cut* (for $\mathcal{A}$) when it has either an admissible connected component or an admissible cut point.

Remarks 1.2.12.
1. In the admissible cut point definition, the stabilizer of A is automatically an allowed edge group since it is a subgroup of G_{e_p}, where e_p is the edge of T corresponding to the vertex p of the Whitehead graph.
2. Since the Whitehead graph has no loop, A contains a vertex.

The following lemma uses that T is locally finite in an essential way.

Lemma 1.2.13 (Dual tree to the Whitehead graph)**.** *Let p a vertex in $\mathrm{Wh}_T(\mathcal{G}, v)$ and W_0 the connected component containing p. If p is a cut point in W_0 (i.e. $W_0 \setminus \{p\}$ is not connected) then p is an admissible cut point of the Whitehead graph.*

Proof. Let p a vertex in $\mathrm{Wh}_T(\mathcal{G}, v)$ whose complement is disconnected. A dual forest to the Whitehead graph can be defined as follows.

First we define the following equivalence relation on the geometric realization of the Whitehead graph W: for all points $x, y \in W$, $x \sim y$ if for all $q \in G_v \cdot p$, x and y are in the same connected component of $W \setminus \{q\}$. Equivalence classes of this relation define a partition of $W \setminus G_v \cdot p$. An equivalence class may contain no vertex of W and that is why we work with the geometric realization of the graph. This partition is coarser than the partition into connected components of $W \setminus G_v \cdot p$.

Then we define the bipartite graph B as follows. There is a vertex u_q for every vertex $q \in G_v \cdot p$. There is also a vertex v_P for every equivalence class P of the equivalence relation defined above. We put an edge between u_q and v_P if $q \in \bar{P}$.

The graph B obtained is a forest because every vertex u_q disconnects all its neighbours in B. It is connected if and only if the Whitehead graph is.

Let W_0 be the connected component of W containing p. Suppose p is a cut point of W_0. The component B_0 of B containing u_p is a finite tree so it has a terminal vertex w. Vertices u_q cannot be terminal since q is a cut vertex so $w = v_P$ for some equivalence class P, which is attached to a vertex $u_q = h \cdot u_p$ of B_0.

The equivalence class P is a connected component of $W_0 \setminus \{h \cdot p\}$ because $G_v \cdot u_p \cap \bar{P} = \{u_q\}$. Taking $A = P$ in the definition, $h \cdot p$ (and thus p) is an admissible cut point. □

We can now state the main theorem of this section. Its proof is given in subsection 1.2.4.

Theorem 1.2.14. *Let $\mathcal{G} \in G$ be a finite collection of loxodromic elements. If $\mathcal{G}$ is simple with respect to $\mathcal{D}^{\mathcal{A}}$ then for all $T \in \mathcal{D}^{\mathcal{A}}$ there exists a vertex $v \in T$ such that $\mathrm{Wh}_T(\mathcal{G}, v)$ has an admissible cut for $\mathcal{A}$.*

Remark 1.2.15. A solvable Baumslag-Solitar group $\mathrm{BS}(1,n) := \langle a, t | tat^{-1} = a^n \rangle$ has no proper special factor (with respect to $\mathcal{D} = \mathcal{D}^{\mathcal{A}_{\min}}$).

Remark 1.2.16. Let T a G-tree and T' a tree obtained from T by subdividing an edge. Then there exists a vertex $v \in T$ such that $\mathrm{Wh}_T(\mathcal{G}, v)$ has an admissible cut if and only if there exists such a vertex in T'.

Indeed T' inherits the Whitehead graphs of T in addition with another Whitehead graph coming from the additional vertex v'. The latter is a graph containing exactly two vertices. If they are joined by an edge, there is no admissible cut. If not, $\mathrm{Wh}_{T'}(\mathcal{G}, v')$ is disconnected, which means for any $g \in \mathcal{G}$, no translate of the axis of g crosses the subdivided edge. In T this edge must then appear as an isolated vertex in the Whitehead graph of one of its endpoints, so some Whitehead graph in T has an admissible cut.

1.2.3 Unfolding lemma

Let $T \in \mathcal{D}^{\mathcal{A}}$. According to remark 1.2.16 we may assume the following : up to performing a finite number of edge subdivisions at the beginning, T has no edge with both ends in the same orbit, i.e. T/G has no loop. This allows us to deal with fewer cases in the proof. The proof is similar to that of an analogous result concerning the case of free products in [GH19, Proposition 5.1]. We will need the following lemma which enables us to perform unfoldings on T or expansions when we find a Whitehead graph with an admissible cut. We allow vertices of valence 2 in the trees considered.

Lemma 1.2.17. *Suppose T/G has no loop. The following conditions are equivalent :*

1. *There exists a Whitehead graph $\mathrm{Wh}_T(\mathcal{G}, v)$ with an admissible cut with respect to $\mathcal{D}^{\mathcal{A}}$.*
2. *There exists a tree $S \in \mathcal{D}^{\mathcal{A}}$ and a non-injective G-equivariant application $f : S \to T$ sending edge to edge or edge to vertex such that for every $g \in \mathcal{G}$, $\|g\|_S = \|g\|_T$.*

In the second condition S/G has also no loop. Moreover the map f can be chosen to be either a fold or a collapse.

Let us start with a preliminary result about lifting the axis of an element $g \in \mathcal{G}$ when performing an unfolding.

Lemma 1.2.18. *Let g be a loxodromic element in G. Let $T, S \in \mathcal{D}^{\mathcal{A}}$ and $f : S \to T$ a simplicial map such that for all edge $e \in T$, edges in the pre-image $f^{-1}(e) := \{\tilde{e} \in E(S)/f(\tilde{e}) = e\}$ all share a common vertex. Suppose that every turn in the axis of g lifts to S, that is to say: for every turn $\{e_1, e_2\}$ in $\mathrm{Axis}_T(g)$ there exists a turn $\{\tilde{e}_1, \tilde{e}_2\}$ in S such that $f(\tilde{e}_1) = e_1, f(\tilde{e}_2) = e_2$.*

Then $\mathrm{Axis}_T(g)$ lifts isometrically in S. Equivalently f is isometric on $\mathrm{Axis}_S(g)$.

Remark 1.2.19. Remark 1.1.24 states that this lemma applies to folds of type A, B and C.

Proof. First of all, given an orientation of an edge $e \in T$, the edges in $f^{-1}(e)$ get a compatible orientation. We will call a set of edges with a common vertex a *star*. If $f^{-1}(e)$ is a star then this orientation is either centripetal or centrifugal.

We claim that for every edge $e \in \mathrm{Axis}_T(g)$ the intersection $f^{-1}(e) \cap \mathrm{Axis}_S(g)$ consists in a unique edge $\tilde{e}$. Moreover, if e, e' are adjacent in T then $\tilde{e}, \tilde{e}'$ are adjacent in S. This yield a continuous application $\mathrm{Axis}_T(g) \to \mathrm{Axis}_S(g)$ which is an inverse for f on the axis. This proves the lemma.

Now let us prove the claim. Let e_1, e_2, e_3 be three consecutive edges in $\mathrm{Axis}_T(g)$ with $t(e_1) = o(e_2)$ and $t(e_2) = o(e_3)$. We will show that $f^{-1}(e_2) \cap \mathrm{Axis}_S(g)$ consists in exactly one edge.

Let A_1, A_2, A_3 be the respective pre-images of e_1, e_2, e_3. They are stars. We endow them with an orientation, either centrifugal or centripetal, compatible with the orientation of their image. Because of orientations in T and since the turns lift, the star A_1 is attached to an end of A_2 if the latter is centripetal, and to the centre if it is centrifugal (see figure 1.8 for a picture of the different cases). On the contrary, A_3 is attached to the centre of A_2 if A_2 is centripetal and to an end if it is centrifugal. In both cases, distance between A_1 and A_3 is 1. There is a unique edge in A_2 which is adjacent to both A_1 and A_3 and we call it $\tilde{e}_2$.

Since $\mathrm{Axis}_T(g) \subset f(\mathrm{Axis}_S(g))$, $\mathrm{Axis}_S(g)$ intersects both A_1 and A_3. Since it is a geodesic, its intersection with A_2 must be the single edge $\tilde{e}_2$. This proves the first part of the claim.

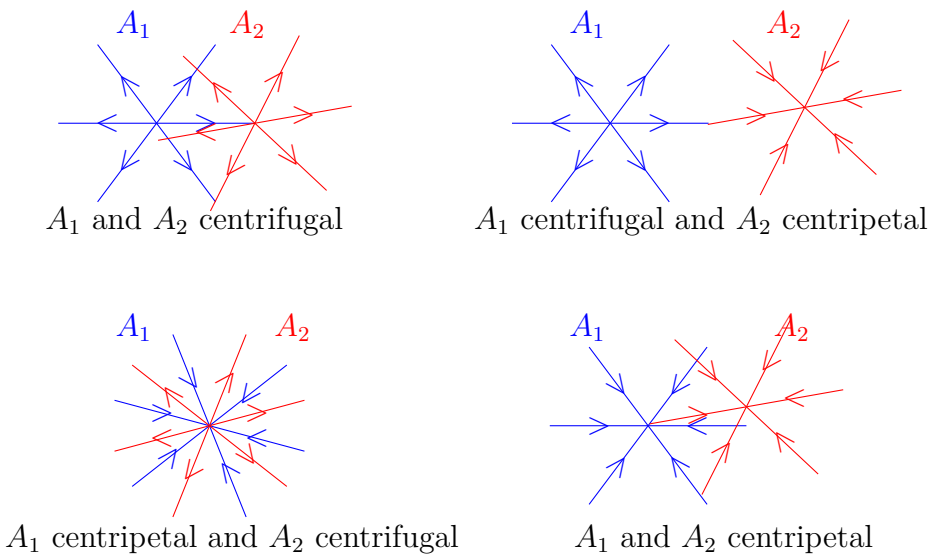

Figure 1.8 – Relative dispositions of the stars depending on their orientations

The second part follows: the lift $\tilde{e}_1$ is adjacent to A_2 and the intersection of two stars is a single point, so it is adjacent to $\tilde{e}_2$. □

Proof of lemma 1.2.17. Suppose that the first condition is true: there is a vertex $v \in T$ such that $\mathrm{Wh}_T(\mathcal{G}, v)$ has an admissible cut. We distinguish several cases based on the shape of $\mathrm{Wh}_T(\mathcal{G}, v)$ and give an application $f : S \to T$ for each case (see figure 1.9). The map f will be a collapse in Case 1, a type A or type C fold in Case 2, and a type B fold in Case 3. Types of folds were defined in subsection 1.1.2. Note that although type A and B folds may look similar, they lead to very different Whitehead graphs which need to be dealt with separately.

Case 1 : Suppose that the Whitehead graph $\mathrm{Wh}_T(\mathcal{G}, v)$ is disconnected and that the stabilizer of a connected component is an allowed edge group of $\mathcal{D}$. In this case, denote the connected components by $C_1, \ldots, C_n$ (note that two connected components might belong to the same orbit). Suppose that $\mathrm{Stab}(C_1)$ is an allowed edge group.

Since C_1 is a connected component of the Whitehead graph, we have $h \cdot C_1 \cap C_1 \neq \varnothing \Rightarrow h \in \mathrm{Stab}(C_1)$. Let $S = T^{\mathrm{Stab}(C_1), C_1}$ obtained by expansion, according to Lemma 1.1.22 with $f : S \to T$ the collapse map. We have $S \in \mathcal{D}^{\mathcal{A}}$ because $\mathrm{Stab}\, C_1 \in \mathcal{A}$ and $T \in \mathcal{D}^{\mathcal{A}}$.

Let $g \in \mathcal{G}$. As no translate of the axis of g crosses a turn between any pair of distinct connected components C_i and C_j, no translate of $\mathrm{Axis}_S(g)$ can cross the added edges, so

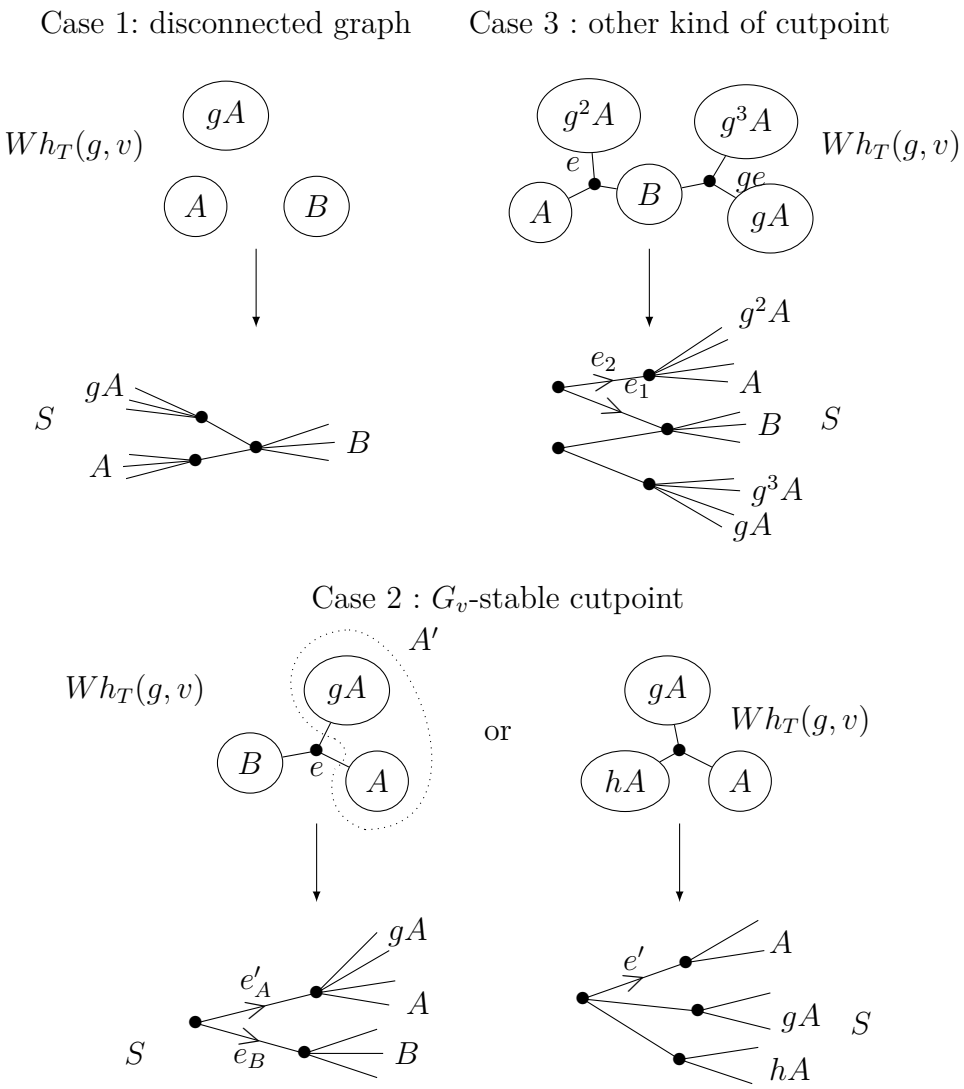

Figure 1.9 – The three cases in the first part of the proof of the lemma. Above, the shape of the Whitehead graph; below, the shape of the tree around the corresponding vertex after transformation.

all turns of $\mathrm{Axis}_T(g)$ lift to S. By lemma 1.2.18 the collapse is isometric on the axes so $\|g\|_S = \|g\|_T$.

In other cases, suppose that the Whitehead graph has no admissible connected component but has a cut point (necessarily admissible, according to lemma 1.2.13). In that case the orbit of the cut point under G_v cannot be the whole graph. The reason is that the graph is finite and without simple loop, and thus cannot have only cut points as vertices.

Case 2: Suppose that the Whitehead graph W has a G_v-invariant admissible cut point. Then G_v fixes an edge in T. Call e such a cut point. Denote by A a connected component of $W \setminus \{e\}$: its stabilizer is a subgroup of G_v so it lies in $\mathcal{A}$.

If $G_v \cdot A \neq W \setminus \{e\}$ denote by B the complement of $G_v \cdot A$ in $W \setminus \{e\}$. The part B is stable under the action of G_v. The subset $A' := G_v \cdot A$ is stable as well. Neither A' nor B are empty. If $\{f, f'\}$ is a turn of a translate of the axis of some $g \in \mathcal{G}$, then $\{f, f'\}$ is included in $G_v \cdot A \cup \{e\}$ or in $B \cup \{e\}$.

We define a new tree S as follows (see Figure 1.9). First we expand T at the vertex v by unattaching edges of A', attaching an edge $e_{A'}$ to v and re-attaching the edges of A' to the other end of $e_{A'}$, which gives the expanded tree $T_1 := T^{G_v, A'}$ (see lemma 1.1.22 for notations). Similarly we unattach edges of B to re-attach them on a new edge e_B with origin v, which gives the tree $T_2 := T_1^{G_v, B}$. The lemma guarantees that T_2 belongs to $\mathcal{D}^{\mathcal{A}}$ since the stabilizer of the new edge is in $\mathcal{A}$. Finally we collapse the edge e of T_2, which is a collapsible edge since its stabilizer is G_v and its ends are not in the same orbit. Let S be the resulting tree. It belongs to $\mathcal{D}^{\mathcal{A}}$. Folding $e_{A'}$ with e_B is a type A fold and yields the original tree T.

Let g be an element of $\mathcal{G}$. Let us prove $\|g\|_S = \|g\|_T$. The pre-image of an edge by the fold $S \to T$ is a star. Every turn in $\mathrm{Axis}_T(g)$ lifts in S: the only turns which do not lift are those of the sort $h \cdot \{A', B\}$. Yet such turns are never crossed by the axis by assumption on the Whitehead graph. According to lemma 1.2.18 f is isometric in restriction to $\mathrm{Axis}_S(g)$ so the translation length of g is the same in T as in S.

On the contrary, if $G \cdot A = W \setminus \{e\}$ then as A is not the only connected component there exists $h \neq 1$ such that $h \cdot A \cap A = \varnothing$. The stabilizer of A is then a subgroup $H \subsetneq G_v$ and H, A satisfy the conditions of the expansion lemma 1.1.22 with H allowed as an edge stabilizer. We perform an expansion at the vertex v as follows: for $u \in G_v/H$, unattach uA and re-attach it on a new edge ue' with origin v. This yields the expanded tree $T^{H,A}$

which belongs to $\mathcal{D}^{\mathcal{A}}$. We then get S by collapsing e. Since e is collapsible we have $S \in \mathcal{D}^{\mathcal{A}}$. When one folds the edges ue' with $u \in G_v/H$, one gets T; the fold is of type C.

Again, for every $g \in \mathcal{G}$, all turns in $\mathrm{Axis}_T(g)$ lift to S so lemma 1.2.18 guarantees $\|g\|_S = \|g\|_T$.

Case 3 : Suppose that W does not have any G_v-invariant cut point. Denote by e a cut point and by W_0 the connected component of W containing e. Let A be a connected component of $W_0 \setminus \{e\}$ which does not contain any element of $G_v \cdot e$. Such a component exists by definition of an admissible cut point. By remark 1.2.12 A contains a vertex. Remember that its stabilizer is an allowed edge group (remark 1.2.12) and is a subgroup of G_e.

Denote by B the complement of $G_v \cdot A$ in $W \setminus G_v \cdot \{e\}$; B may be empty and may intersect $W \setminus W_0$. Again B is stable by G_v. We also define $A' := G_e \cdot A$.

Since $G_v \cdot e$ has at least two elements, $\{e\} \cup A'$ is a proper subset of $\mathcal{E}_v$, so even when B is empty, we may use lemma 1.1.22 to do the following expansions.

See figure 1.10 for a closer illustration of the case. First we do an expansion at vertex v: we partition the set of edges into $B \sqcup \bigsqcup_{h \in G_v/G_e}(\{e\} \cup A')$. We get the tree $T_1 := T^{G_e, \{e\} \cup A'}$ with notations of lemma 1.1.22: we replace the vertex v by a star with $|G_v/G_e|$ branches. Then we attach B to the centre of the star, and edges in $h \cdot (\{e\} \cup A')$ to $h \cdot G_e$. We call e_1 the edge joining $\{e\} \cup A'$ to the centre of the star (which we will still call v). We call w the origin of e_1.

Then we perform a second expansion at w which is the origin of e_1, $\bar{e}$ and of the edges of A' and has stabilizer G_e. The tree $T_2 := T_1^{G_e, A'}$ may be described as follows: we unattach the edges of A' then re-attach them on a new edge e_2 with origin w (see figure 1.10).

Finally consider the collapse $S = T_2/\sim_e$. Since $G_e \in \mathcal{A}$, T_1, T_2 and S are in $\mathcal{D}^{\mathcal{A}}$. There is an application $S \to T$ which sends e_1 and e_2 on e and it is the (type B) fold of e_1 with e_2.

Let g be an element of $\mathcal{G}$. The only turns of T at v which may be crossed by translates of the axis of g are those of the sort $h \cdot \{A, A\}$, $h \cdot \{A, e\}$, $\{B, B\}$, $h \cdot \{B, e\}$, $h \cdot \{e, h' \cdot e\}$ with $h, h' \in G_v$. All these turns lift to S. According to lemma 1.2.18, the whole axis lifts isometrically so $\|g\|_T = \|g\|_S$.

We have proved the existence of S and f in all cases where the Whitehead graph has an admissible cut.

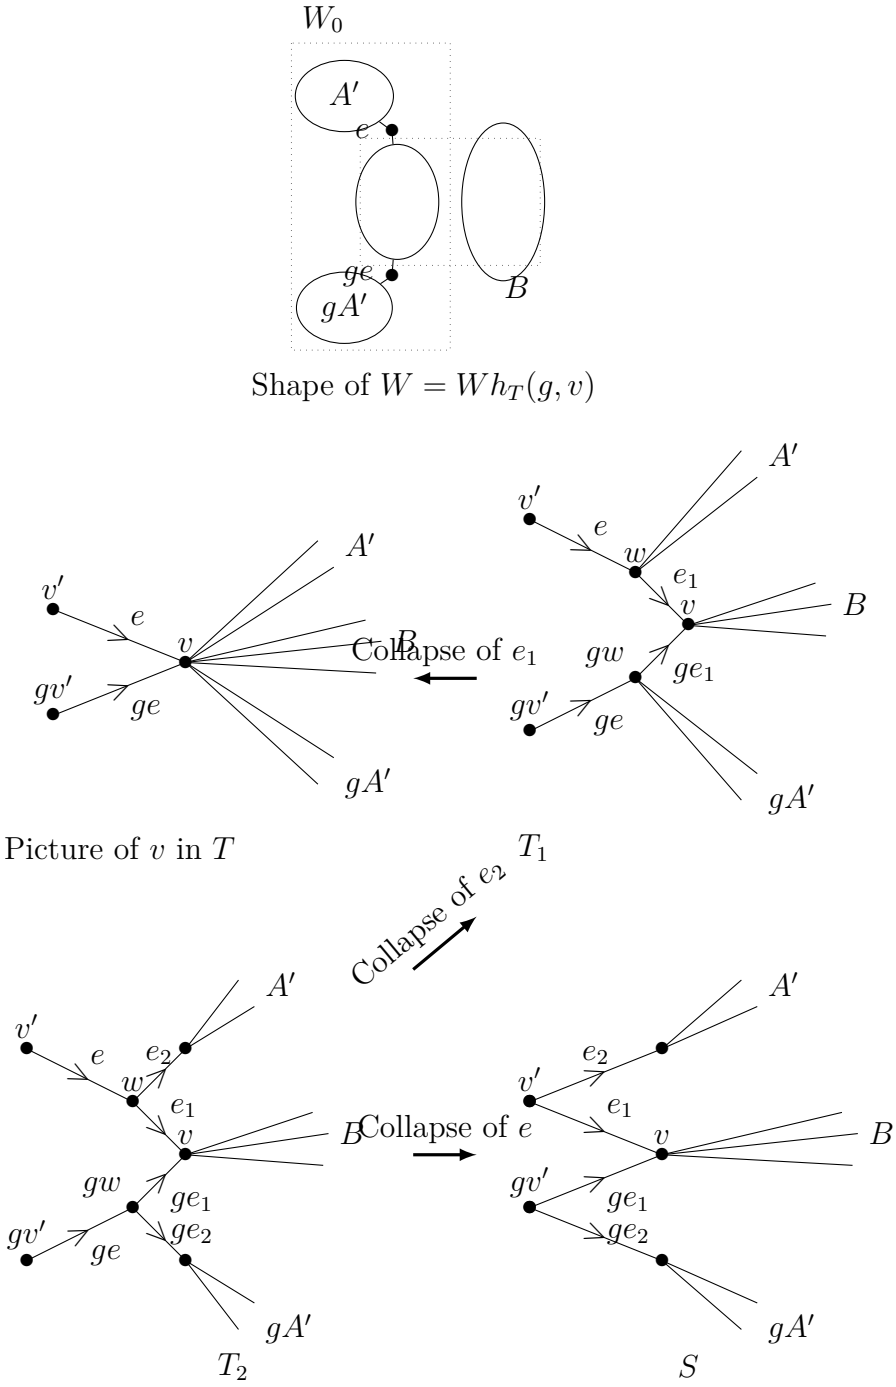

Figure 1.10 – Steps of the unfolding in the case 3, where the orbit of the cut point has several elements. Arrows represent collapses.

Conversely, suppose there exists S and $f : S \to T$ non-injective, sending edge to edge or edge to vertex, such that $\|g\|_S = \|g\|_T$ for every $g \in \mathcal{G}$. According to [BF91], this application may be considered as a composition of collapses and folds.

Let us consider only the last collapse or last fold. Since neither folds nor collapses can increase translation length, this application satisfies the second condition of the lemma. We may then suppose f is either a fold or a collapse, which simplifies the proof. In both cases f is 1-Lipschitz. The assumption about translation lengths implies that for every $g \in \mathcal{G}$, f is isometric on the axis of g. Therefore all turns in $\mathrm{Axis}_T(g)$ lift to S.

If f is a collapse, as it does not change deformation space, it is a quasi-isometry, so connected components of the subforest collapsed by f are bounded.

Let $v \in T$ be such that the subtree $f^{-1}(v)$ is not reduced to a point. As f is a collapse, the pre-image of any edge in T is a single edge in S. In the Whitehead graph $\mathrm{Wh}_T(\mathcal{G}, v)$, any two vertices joined by an edge correspond to edges of T in the same connected component of $\overline{S \setminus f^{-1}(\{v\})}$. Otherwise $\mathrm{Axis}_S(g)$ contains an edge collapsed by f. Therefore the Whitehead graph has at least as many connected components as $\overline{S \setminus f^{-1}(\{v\})}$ which is not connected.

Let e a collapsed edge in S whose image is v. As e is collapsible, it has an end w such that $G_w = G_e$ and such that w is terminal in $f^{-1}(v)$. The vertex w belongs to the boundary of $S \setminus f^{-1}(v)$ since S has no valence 1 vertex. Every connected component of $S \setminus f^{-1}(\{v\})$ whose boundary is w has stabilizer included in G_e. The stabilizers of all corresponding components in $\mathrm{Wh}_T(\mathcal{G}, v)$ are then subgroups of G_e, so they are allowed edge groups. Thus the Whitehead graph has an admissible cut.

If f is a fold, it is defined by two edges of S with same origin w: call them e_1 and e_2. Call their endpoints v_1 and v_2. Call e' the edge of T which is the image of e_1 and e_2, w' its initial vertex and v' its terminal vertex (which is the image of v_1 and v_2). The vertices v' and w' are in different orbits as we supposed that T/G is without simple loop.

We will prove that $\mathrm{Wh}_T(\mathcal{G}, v') \setminus \{e'\}$ is disconnected and the stabilizer of at least one of its connected components E_1 is in $\mathcal{A}$. Lemma 1.2.13 states that this implies that $\mathrm{Wh}_T(\mathcal{G}, v')$ has an admissible cut.

Figure 1.12 recaps all different cases of folds and associated shapes of graphs.

Three different kind of folds may occur, which correspond to cases A, B and C.

1. $G_{e_i} = G_{v_i}$ for $i \in \{1, 2\}$ and e_1, e_2 lie in different orbits (type A)
2. $G_{v_1} \subset G_{e_2}$, up to permutation of indices, and e_1, e_2 lie in different orbits (type B)

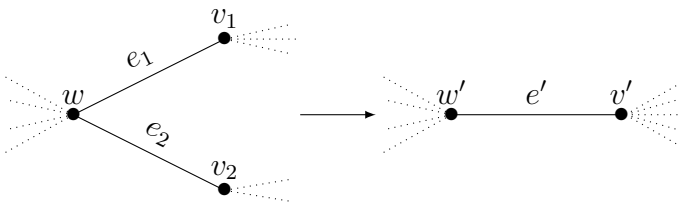

Figure 1.11 – A fold as described.

3. $e_2 = he_1$ for some $h \in G_v$ and $G_{e_i} = G_{v_i}$ (type C)

In the three cases, define $\tilde{E}_1$ to be the set of edges of S with origin v_1, except $\bar{e}_1$. Let E_1 be the image of $\tilde{E}_1$ in T. Define E_2 similarly.

Since all turns represented in $\mathrm{Wh}_T(\mathcal{G}, v')$ lift to S, any edge of $\mathrm{Wh}_T(\mathcal{G}, v')$ with one endpoint in E_1 joins E_1 to itself or to $\bar{e}'$. In particular, E_1 and E_2 are in distinct connected components of $\mathrm{Wh}_T(\mathcal{G}, v') \setminus \{e'\}$. Therefore $\mathrm{Wh}_T(\mathcal{G}, v')$ has an admissible cut. □

1.2.4 Proof of the theorem

Lemma 1.2.20. *(Expansion of non-cyclic vertex groups) Let $R \in \bar{\mathcal{D}}^{\mathcal{A}}$ whose certain vertices $v_1, \dots, v_k$ (in different orbits) have stabilizers $H_1, \dots, H_k$ some special factors. Let $T \in \mathcal{D}$ and T_i the minimal subtree for H_i in T. There exists a G-tree $S \in \mathcal{D}^{\mathcal{A}}$ and a map $f : S \to T$ and a collapse $\pi : S \to R$ such that:*

— *the image by π of the collapsed subforest is $G \cdot \{v_1, \dots, v_k\}$,*

— *for all $i \in \{1, \dots, k\}$, $f_{|\pi^{-1}(v_i)}$ is an isomorphism to T_i.*

In other words, it is possible to blow up R by replacing v_i by T_i.

A proof of this result is given in [GL17, Proposition 2.2]. The key assumption is the fact that all edge groups in R are elliptic in T so we can attach edges of R to the subtrees which replace the vertices. Here we suppose that R is the result of the collapse of some $\tilde{R} \in \mathcal{D}$, so its edge groups are also edge groups in $\tilde{R}$ and are elliptic in any tree in $\mathcal{D}^{\mathcal{A}}$.

Proof of theorem 1.2.14.. Let G be a GBS group and let $\mathcal{G} := \{g_1, \dots, g_k\}$ be a finite collection of loxodromic elements of G. Let $T \in \mathcal{D}^{\mathcal{A}}$. Suppose that $\mathcal{G}$ is simple with respect to $\mathcal{D}^{\mathcal{A}}$, that is, there exists a non-trivial G-tree $R \in \bar{\mathcal{D}}^{\mathcal{A}}$ such that every $g_i \in \mathcal{G}$ fixes a vertex $v_i \in R$.

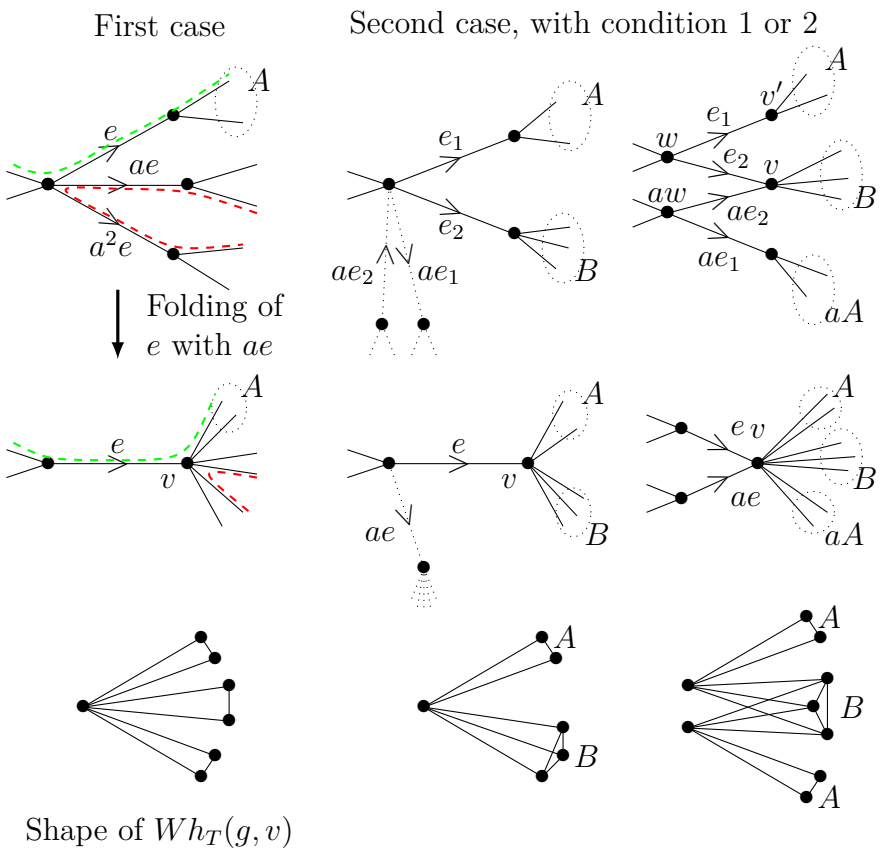

Figure 1.12 – Summary of all different cases of folds which may occur and corresponding shape of associated Whitehead graph

For every $1 \leq i \leq k$, let T_i be the minimal G_{v_i}-invariant subtree of T. We obtain a new tree $S \in \mathcal{D}^{\mathcal{A}}$ by applying lemma 1.2.20 starting from R so that for every $i \in \{1, \dots, k\}$, the vertex v_i is replaced by a copy of T_i.

Let $f : S \to T$ be the map given by Lemma 1.2.17. If f is injective then it is an isomorphism (surjectivity is obtained by minimality of the image for the action of G). In that case, for every $i \in \{1, \dots, k\}$, the axis of g_i avoids images of edges in S coming from edges in R. Therefore some Whitehead graph at an end of such an edge has an isolated vertex and has an admissible cut.

When f is not injective, we showed Lemma 1.2.17 that there is a vertex at which the Whitehead graph has an admissible cut, which proves the theorem. □

1.3 Algorithm

This part is dedicated to a theorem which states that one can decide algorithmically whether a collection of loxodromic elements of G is simple with respect to some $\mathcal{D}^{\mathcal{A}}$, or not.

Before stating the theorem, let us explain how we deal with the set $\mathcal{A}$ of allowed edge groups. Let Γ be a graph of groups representing G; it has finitely many vertices $v_1, \dots, v_n$.

Definition 1.3.1. For H a subgroup of G we denote by $\mathcal{A}_H$ the family of subgroups $\{A \cap H, A \in \mathcal{A}\}$. Since $\mathcal{A}$ is stable by taking subgroups, $\mathcal{A}_H$ is a subfamily of $\mathcal{A}$.

For each $v \in V(\Gamma)$ let I_v be a family of positive integers. We say that $(I_v)_{v \in V(\Gamma)}$ *represents* $\mathcal{A}$ if for every $v \in V$, $\mathcal{A}_{|G_v}$ is the set of all subgroups of G_v whose index is a multiple of an element in I_v.

Examples 1.3.2. — $I_v = \{1\}$ for each $v \in \Gamma$ represents $\mathcal{A}_{\mathbb{Z}}$.

— Suppose Γ is reduced. Let I_v be the set of absolute values of labels at v, then $(I_v)_{v \in \Gamma}$ represents $\mathcal{A}_{\min}$.

— Given Γ a graph of groups for G, we can choose $\mathcal{A}$ to be the set of bi-elliptic subgroups in Γ. If so, $\mathcal{A}$ is represented by $(I_v)_{v \in V(\Gamma)}$ where I_v is the set of all labels at v.

In the sequel we assume that $\mathcal{A}$ is represented by a family $(I_v)_{v \in V(\Gamma)}$ of finite sets.

Here is the theorem which we will prove.

Theorem 1.3.3. *There is an algorithm which takes as input:*

— *a graph of cyclic groups* Γ_G *representing* G

— *a marking* $G \to \pi_1(\Gamma_G, v_0)$

— *a finite family of finite subsets* $(I_v)_{v \in V(\Gamma_G)}$ *representing a family* $\mathcal{A}$ *of subgroups of* G

— *a finite collection of loxodromic elements* $\mathcal{G} \subset G$.

which decides whether there exists a system of proper special factors $\mathcal{H}$ *of* G *with respect to* $\mathcal{D}^{\mathcal{A}}$ *such that* $\mathcal{G} \preceq \mathcal{H}$*, and then returns such a system when it exists.*

We show that the construction of Whitehead graphs is algorithmic, and that the construction of the map of Lemma 1.2.17 is algorithmic too. Finally, after giving a description of the algorithm, we prove that it terminates.

1.3.1 Paths in graphs of groups, fundamental group

In this section we give some general definitions and results about GBS groups. For algorithmic purposes it is more convenient to work with graphs of groups than with trees.

Let G be a GBS group. We suppose it is given by a graph of cyclic groups Γ_G. From this graph we deduce a presentation $\langle s_1, \dots, s_n | r_1, \dots, r_m \rangle$. We will use this presentation to define a marking in other graphs of groups below.

Let Γ be any graph of cyclic groups. It is given as a set of vertices and edges, a label $\lambda(e)$ for each edge e and a set of generators $(a_v)_{v \in V(\Gamma)}$ for all vertex groups.

We define $\pi_1(\Gamma)$ as a subgroup of the Bass group $B(\Gamma)$ like in [Ser77, Definition (a), part 5.1]. The group $B(\Gamma)$ has the following presentation: generators are

— elements a_v indexed by vertices of Γ, where a_v is to be thought of as a generator of the vertex group $G_v \simeq \mathbb{Z}$

— elements t_e indexed by oriented edges in Γ.

Relations are the following:

1. for all edge $e \in \Gamma$ we have $t_e = t_{\bar{e}}^{-1}$
2. for all $e \in \Gamma$ with initial vertex u and terminal vertex v, with labels $\lambda(e) = p$ and $\lambda(\bar{e}) = q$, we have $a_v^q = t_e a_u^p t_e^{-1}$

A *path* α *in the graph of groups* is a pair $\alpha = (w, \gamma)$ where

— γ is a path $v_0, e_1, \dots, e_n, v_n$ in the underlying graph of Γ, where the v_i are vertices and e_i are edges such that $t(e_i) = v_i = o(e_{i+1})$,

— w is a word $a_0 t_{e_1} a_1 \dots t_{e_n} a_n$ where $a_i \in G_{v_i}$ for all $i \in \{0, \dots, n\}$.

We denote by $[\alpha]$ the element of $B(\Gamma)$ represented by w.

The *length* of a path (w, γ) is the number of edges in γ. Its initial and terminal vertices are v_0 and v_n respectively.

Let $\alpha = (a_0 t_1 \dots a_n; v_0, e_1, \dots, v_n)$ and $\alpha' = (a'_0 t'_1 \dots a'_m; v'_0, e'_1, \dots, v'_m)$ be two paths in Γ such that $v_n = v'_0$. The *concatenation* of α and α' is the path

$$\alpha \cdot \alpha' = (a_0 t_1 \dots t_n b t'_1 a'_1 \dots a'_m; v_0, e_1, \dots, v_n, e'_1, \dots, v'_m)$$

where b is the element of G_{v_n} equal to $a_n a'_0$.

Let $v \in \Gamma$. A *loop based at* v in Γ is a path with initial and terminal vertices equal to v.

To a loop $\alpha = (w, \gamma)$, one associates the corresponding element $[\alpha] := [w] \in B(\Gamma)$.

If α, α' are loops in Γ based at v then $[\alpha \cdot \alpha'] = [\alpha] \cdot [\alpha']$.

Fix a vertex v in Γ. The fundamental group $\pi_1(\Gamma, v)$ is the subgroup of $B(\Gamma)$ consisting of the elements of $B(\Gamma)$ associated to loops based in v.

Remark 1.3.4. If a word in the generators of $B(\Gamma)$ corresponds to a path in the graph, then this loop is unique. We really mean the word as a sequence fo letters and not the corresponding element of $B(\Gamma)$. Thus the word is sufficient to describe a path in the graph of groups, and we will use the word on its own when the description in terms of edges and vertices is not needed.

A *marking* of Γ is a map $\{s_i, 1 \leq i \leq n\} \to \pi_1(\Gamma, v)$ which associates a loop based at v in Γ to each generator of G, such that it induces an isomorphism $G \to \pi_1(\Gamma, v)$.

For every $g \in \mathcal{G}$, given the expression of g as a word in the generators $\{s_i, 1 \leq i \leq n\}$, we can determine a loop in Γ based at v which represents g.

We call a path in Γ given by a word w *reduced* if no subword in w is of the form $t_e a t_{\bar{e}}$ with $a \in i_e(G_e)$. Note that here vertex groups are cyclic so determining if a belongs to $i_e(G_e)$ boils down to a question of divisibility.

We call a loop α *cyclically reduced* if the concatenation $\alpha \cdot \alpha$ is reduced.

We may modify the marking by the following process. Given another vertex v' and a path α in the graph of groups Γ from v to v', there is an isomorphism $\pi_1(\Gamma, v) \simeq \pi_1(\Gamma, v')$ defined by

$$\sigma_\alpha : \pi_1(\Gamma, v) \to \pi_1(\Gamma, v')$$
$$[h] \mapsto [\bar{\alpha} h \alpha]$$

Lemma 1.3.5. *A path $h = (w, \gamma)$ can be reduced algorithmically, i.e. there is an algorithm which finds a reduced path h' such that $[h] = [h']$ in $B(\Gamma)$.*

A loop can be cyclically reduced algorithmically: for any loop α, one can find a cyclically reduced loop α' and a path β such that $[\alpha] = [\bar{\beta}\alpha'\beta]$.

The proof is standard and straightforward. We leave it to the reader.

Lemma 1.3.6. *Suppose that γ is a reduced loop based at v and that $[\gamma] \in G_v$. Then γ has length 0.*

The proof for this fact is in [Ser77, 5.2, Theorem 11].

With the elements above we define the universal cover $T_{\Gamma,v}$ of the graph of groups Γ. It is a graph defined as follows. The set of vertices is

$$\tilde{V} = \{\text{paths in } \Gamma \text{ with initial vertex } v\} / \sim$$

where $\gamma \sim \gamma'$ if γ and γ' have the same terminal vertex $v_i \in \Gamma$ and $[\gamma]^{-1}[\gamma'] \in G_{v_i}$. Denote by $[\gamma]_V$ the vertex associated to the path γ. The group $\pi_1(\Gamma, v)$ acts on $\tilde{V}$ by left concatenation.

Note that checking whether two paths define the same vertex boils down to checking whether a path in Γ can be reduced to a length zero path, by lemma 1.3.6, so it is algorithmic.

Let $\tilde{v} = [1]_V$. It is a lift for the base point.

The oriented edges of (T, v) are defined as follows:

$$\tilde{E} = \left\{ (\alpha, at_e) / \ \alpha \text{ path in } \Gamma \text{ from } v \text{ to } v', e \in \mathcal{E}_{v'}, a \in G_{v'} \right\} / \sim$$

where $\mathcal{E}_{v'}$ is the set of edges with origin v'. The equivalence relation $\sim$ is defined by $(\alpha, at_e) \sim (\alpha', a't_{e'})$ if and only if $e = e'$ and $a^{-1}\alpha^{-1}\alpha' a' \in i_{\bar{e}}(G_e)$. The origin of this edge is $[\alpha]_V$ and its terminus is $[\alpha \cdot at_e]_V$.

Denote by $[(\alpha, at_e)]_E$ the equivalence class of (α, at_e). The group $\pi_1(\Gamma, v)$ acts on $\tilde{E}$ by left concatenation of α. We have the relation $[(\alpha, at_e)]_E = [(\alpha \cdot a, t_e)]_E$ for all $a \in G_{v'}$.

The edge with opposite orientation is $\overline{[(\alpha, at_e)]_E} = [\alpha \cdot at_e, t_{\bar{e}}]$.

The graph defined is a tree. The quotient of (T, v) under the action of $\pi_1(\Gamma, v)$ is Γ; the projection of $[\gamma]_V$ is the last vertex of γ (see [Ser77] for a proof).

Note that for different base points v, v' the universal covers $T_{\Gamma,v}$ and $T_{\Gamma,v'}$ are isomorphic. Given α joining v to v', we define an isomorphism between $T_{\Gamma,v}$ and $T_{\Gamma,v'}$ by $[\gamma]_V \mapsto [\bar{\alpha} \cdot \gamma]_V$ (this is well-defined since it does not depend on the representative γ chosen). A marking $G \to \pi_1(\Gamma, *)$ yields an action of G on $T_{\Gamma,*}$. If we identify $\pi_1(\Gamma, v)$ with $\pi_1(\Gamma, v')$ with the isomorphism σ_α as above, the isomorphism between the trees is G-equivariant.

Lemma 1.3.7. *Let α be a path in Γ with first vertex v. It can be lifted to a path in T with first vertex $\tilde{v}$ and last vertex p_α, the equivalence class of α in $\tilde{V}$.*

Suppose α is reduced. Then the distance between $\tilde{v}$ and p_α in (T, v) is equal to the length of α.

Proof. We follow the proof of [Bee15, Proposition 2.6].

First of all we construct the path between $\tilde{v}$ and p_α. Denote by n the length of α. Write (as a word) $\alpha = a_0 t_{e_1} a_1 \dots a_n$, and $\alpha_i := a_0 t_{e_1} a_1 \dots a_i$ for $i \leq n$.

Let $\tilde{v}_i = [\alpha_i]_V$ and $\tilde{e}_i = [(\alpha_{i-1}, t_{e_i})]_E$.

The path $\tilde{v}, \tilde{e}_1, \tilde{v}_1, \dots, \tilde{e}_n, \tilde{v}_n = g \cdot \tilde{v}$ is a path in T. Let us show that it is a geodesic, that is, no consecutive edge are opposite.

Suppose $\overline{\tilde{e}_{i+1}} = \tilde{e}_i$ for some $i \in \{0, \dots, n\}$. Then

$$\overline{[(\alpha_i, t_{e_{i+1}})]_E} = [(\alpha_i \cdot t_{e_{i+1}}, \bar{t}_{e_{i+1}})]_E = [(\alpha_{i-1}, t_{e_i})]_E$$

so $\bar{e}_i = e_{i+1}$ and $t_{e_i} a_i \bar{t}_{e_i} \in i_{e_i}(G_{e_i})$. Since the word is reduced this leads to a contradiction. Thus the lift of the path α is a geodesic in T. □

Lemma 1.3.8. *Let Γ be a labelled graph and g be an element in $\pi_1(\Gamma, v)$ represented by a cyclically reduced loop γ. Let T be the universal cover of Γ defined as above.*

If the loop has length zero, then g is elliptic. Otherwise g is loxodromic and the path in T defined by γ, joining $\tilde{v}$ to $g\tilde{v}$, is a fundamental domain of the axis of g in T.

In particular, for each loxodromic element $g \in G$, one can compute a fundamental domain of its axis.

Proof. Let γ be a cyclically reduced loop in Γ such that $[\gamma] = g$. Let $\tilde{v}$ be the lift of the base point v in T. Since γ is cyclically reduced, lemma 1.3.7 ensures that $[\tilde{v}, g^2 \cdot \tilde{v}]$ is a geodesic. Therefore $\tilde{v}$ belongs to the axis of g in T. As a result the path $[\tilde{v}, g \cdot \tilde{v}]$ is a fundamental domain of the axis of g. □

1.3.2 Algorithmicity of Whitehead graph computation and unfoldings

Lemma 1.3.9. *Let Γ be a marked graph of groups for G. Let T be its universal cover. Let x be a vertex in T represented as $x = [\gamma_x]_V$ for some path γ in Γ. Let G_x be the stabilizer of x, not to be confused with vertex groups in Γ. The subgroup G_x is generated by $a_x := \gamma_x a_{\pi(x)} \bar{\gamma}_x$ and one can compute*

- *the link* $\mathrm{lk}(x)$
- *the action of* a_x *on* $\mathrm{lk}(x)$.

Proof. Since T is a tree, $\mathrm{lk}(x)$ is the collection of edges of T with initial vertex x. Denote by π the quotient map $T \to \Gamma$. Any edge of T with origin x has a unique representative of the form $(\gamma_x, a^k_{\pi(x)} t_e)$ with $e \in \mathrm{lk}(\pi(x))$ and $0 \leq k < \lambda(e)$.

All such edges can be listed algorithmically since the indices of edge groups in vertex groups are all finite.

For every $y := [\gamma_x, a^k_{\pi(x)} t_e]_E \in \mathrm{lk}(x)$ we have $a_x \cdot y = [\gamma_x, a^{k+1 \bmod \lambda(e)}_{\pi(x)} t_e]_E$. □

Let $g \in G$ be a loxodromic element. With the input of a graph of groups and a loop for g we can compute the Whitehead graphs $\mathrm{Wh}_T(\{g\}, v)$:

Lemma 1.3.10. *Let Γ be a graph of cyclic groups. Let (w, γ) be a cyclically reduced loop in Γ based in v representing some $g \in \pi_1(\Gamma)$. Let T be the universal cover of Γ at basepoint v and let x be a vertex in T. The computation of the Whitehead graphs $\mathrm{Wh}_T(\{g\}, x)$ is algorithmic.*

Proof. By lemma 1.3.9 we can compute the link of x and the action of G_x on it.

By lemma 1.3.8 we may compute a fundamental domain of the axis of g^2 (or equivalently a pair of consecutive fundamental domains of g). All orbits of turns crossed by the axis of g appear in this segment. For every turn τ of $\mathrm{Axis}(g)$ based at a point in the orbit of x we can find a pair of edges in $\mathrm{lk}(x)$ forming a turn in the same orbit, and using the action of G_x, we can find all turns at x in the orbit of τ.

The elements of $\mathrm{lk}(x)$ form the vertices of $\mathrm{Wh}_T(g,x)$ and the turns computed above are edges. □

Corollary 1.3.11. *Let $\mathcal{G} \subset G$ be a finite collection of loxodromic elements of G. Then the computation of the Whitehead graphs $\mathrm{Wh}_T(\mathcal{G}, v)$ is algorithmic.* □

Lemma 1.3.12. *Given a Whitehead graph $\mathrm{Wh}_T(\mathcal{G}, v)$, the action of G_v on it and the set I_v, one can decide algorithmically whether it has an admissible cut.*

Proof. Finding connected components in a finite graph is algorithmic. The set I_v allows to check the admissibility of connected components by calculating stabilizers. □

Lemma 1.2.17 states that given a Whitehead graph with an admissible cut there exists a non-injective map $f : S \to T$ which preserves translation length of every $g \in \mathcal{G}$. In the proof we actually gave a construction of such a map. We are going to prove that this construction can be done algorithmically. First we show that collapses and expansions can be done algorithmically:

Lemma 1.3.13. *There is an algorithm which takes as input a graph of groups Γ, a collapsible edge $\epsilon = vw$ of Γ such that $G_w = G_\epsilon$, a marking $\psi : G \to \pi_1(\Gamma, v)$ and outputs*

- *a marked graph of groups (Γ', ψ') whose universal cover is the tree T' obtained by the collapse of the orbit ϵ of the universal cover T of (Γ, ψ)*
- *an isomorphism $\phi : \pi_1(\Gamma, v) \to \pi_1(\Gamma', v')$ such that $\phi \circ \psi = \psi'$.*

Proof. First construct the graph of groups Γ': we collapse the collapsible edge ϵ in Γ. Since $G_w = G_\epsilon$ we have $\lambda(\bar{\epsilon}) = \pm 1$. We create the new graph of group Γ' by deleting ϵ and w and redefining any edge with origin w by attaching it to v instead. The label of such edges is multiplied by $\pm\lambda(\epsilon)$. There is a map between the underlying graphs $f : \Gamma \to \Gamma'$ which sends each vertex to the corresponding vertex in Γ' (which we will write with a $'$) and sends w to v', and sends any edge except ϵ to the corresponding edge and sends e to the vertex v'.

Now let us build a morphism $\phi : B(\Gamma) \to B(\Gamma')$:

$$\begin{aligned} a_u &\mapsto a_{u'} \text{ for } u \neq w \\ a_w &\mapsto a_{v'}^{\pm\lambda(\epsilon)} \\ t_e &\mapsto t_{e'} \text{ for } \epsilon \neq e \\ t_\epsilon &\mapsto 1 \end{aligned}$$

This morphism sends any path in the graph of groups Γ to a path in Γ'. In particular it induces a morphism $\pi_1(\Gamma, v) \to \pi_1(\Gamma', v')$. Moreover the induced morphism is an isomorphism between both fundamental groups. We define $\psi' = \phi \circ \psi$. It is a marking on Γ'.

The morphism ϕ induces a map $(T, v) \to (T', v')$ because it sends paths to paths and preserves the equivalence relations defining trees. This map is a collapse map and it is G-equivariant for the markings ψ and ψ'.

The graph Γ' and the marking ψ' can be computed from Γ and ψ, which proves the lemma. □

We have a similar result for expansions. It is slightly more difficult since while collapses may be defined from data in the quotient, the definition of an expansion requires some information in the tree.

Lemma 1.3.14. *There is an algorithm which takes as input*

- *a marked graph of groups* (Γ, ψ)
- *a vertex* $v \in \Gamma$
- *a subset* $S \subsetneq E(\tilde{v})$*, where* $\tilde{v}$ *is the basepoint of the universal cover* (T, v)*, such that* $\forall g \in G_{\tilde{v}}, \quad gS \cap S \neq \varnothing \Rightarrow g \in \mathrm{Stab}(S)$

and outputs

- *a marked graph* (Γ', ψ') *whose universal cover is* $T' = T^{\mathrm{Stab}(S),S}$ *(see lemma 1.1.22)*
- *an isomorphism* $\phi : \pi_1(\Gamma, v) \to \pi_1(\Gamma', v')$ *such that* $\phi \circ \psi = \psi'$*.*

Proof. Like in the proof of the collapse in lemma 1.3.14, we first construct an oriented graph, then construct a map between $B(\Gamma)$ and $B(\Gamma')$ which induces a map between trees.

Let Γ' be the labelled graph obtained as follows. The vertices of Γ' are the vertices of Γ along with another vertex w'. To distinguish them from the vertices of Γ we write their names with a $'$. The edges are redefined according to the following rule (see figure 1.13). Let $n := [G_{\tilde{v}} : \mathrm{Stab}(S)]$. Let $\pi : T \to \Gamma$ be the quotient map.

1. Add an edge ϵ with origin v' and terminal vertex w' with label n near v' and 1 near w'.
2. For $e \in E(\Gamma)$ not in $\pi(S)$, keep e with same origin and same label $\lambda(e)$.
3. For $e \in E(\Gamma)$ in $\pi(S)$ we have $o(e) = v$. Redefine its origin to be w' and its label to $\lambda(e)/n$.

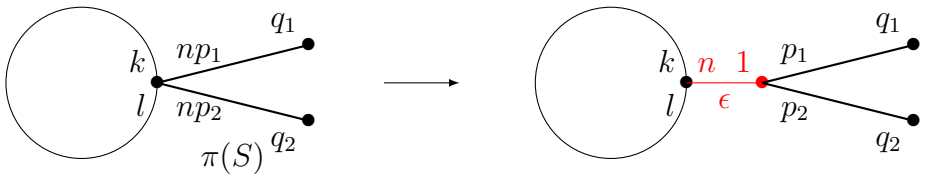

Figure 1.13 – Redefinition of edges

This defines the origin of all oriented edges of Γ'. We then define $t(e) = o(\bar{e})$ for all $e \in E(\Gamma')$.

Let $f : \Gamma' \to \Gamma$ be the natural collapse map.

For each edge e with origin v in Γ, let $\tilde{e} := [(1, t_e)]_E$ be its standard lift in T with origin $\tilde{v}$. If $e \in \pi(S)$ choose $c(e)$ such that $e_S := a_v^{c(e)}\tilde{e} \in S$. Note that $\pi^{-1}(e) \cap S = \langle a_v^n \rangle \tilde{e}_S$.

Let $\chi : B(\Gamma') \to B(\Gamma)$ be the following morphism:

$$\begin{aligned}
a_{u'} &\mapsto a_u \text{ for } u' \neq w' \\
a_{w'} &\mapsto a_v^n \\
t_{e'} &\mapsto t_e \text{ for } e \notin \pi(S) \cup \pi(\bar{S}) \\
t_{e'} &\mapsto a_v^{c(e)} t_e \text{ for } e \in \pi(S) \\
t_\epsilon &\mapsto 1
\end{aligned}$$

and define the morphism $\phi : B(\Gamma) \to B(\Gamma')$:

$$\begin{aligned}
a_u &\mapsto a_{u'} \\
t_e &\mapsto t_{e'} \text{ for } e \notin \pi(S) \cup \pi(\bar{S}) \\
t_e &\mapsto a_v^{-c(e)} t_\epsilon t_{e'} \text{ for } e \in \pi(S)
\end{aligned}$$

We can check that $\chi \circ \phi = \mathrm{id}_{B(\Gamma)}$. Both these morphism send paths to paths so they induce morphisms between fundamental groups. One also easily checks that $\phi \circ \chi_{|\pi_1(\Gamma', v')} = \mathrm{id}_{\pi_1(\Gamma', v')}$. Thus χ induces an isomorphism $\pi_1(\Gamma', v') \to \pi_1(\Gamma, v)$ and the restriction of ϕ is its inverse.

Define the marking $\psi' = \phi \circ \psi$ of Γ'.

The morphism χ sends paths to paths, so it induces a map $\tilde{f} : (T', v') \to (T, v)$ between universal covers. This map is G-equivariant for the markings ψ and ψ'.

We want to prove that $\tilde{f}$ is the collapse map of the orbit ϵ and that $T' = T^{\mathrm{Stab}\, S, S}$.

Let $\tilde{v}' = [1_{v'}]_V$ be the basepoint in T'. Let $\tilde{w}' = [1_{v'} t_\epsilon 1_{w'}]_V$ and $\tilde{\epsilon} = [(1_{v'}, t_\epsilon)]_E$. In order to prove that $T' = T^{\mathrm{Stab}(S),S}$ we need to show that the link of the vertex $\tilde{w}'$ is the pre-image of S by $\tilde{f}$, together with the edge $\bar{\tilde{\epsilon}}$.

Note that $\tilde{f}$ collapses the orbit ϵ. To prove that it is a collapse map, we also need the following fact: $\tilde{f}$ is injective on the set of edges not in the orbit of ϵ. Indeed, for $e \in E(T') \setminus G \cdot \epsilon$, $\phi \circ \chi(G_e) = G_e$. Thus two distinct edges in the same orbit cannot be sent to the same edge.

Note that edges in different orbits are also sent to images in different orbits. This proves that $\tilde{f}$ is injective on the set of edges. Thus it is injective on the interior of $T' \setminus G \cdot \tilde{\epsilon}$ and it is the collapse map associated to the orbit ϵ.

Let $S' = \mathrm{lk}(\tilde{w}) \setminus \{\tilde{\epsilon}\}$. Let us show that it is contained in the preimage of S. Let $\tilde{e}' := [(1_{v'} t_\epsilon a_{w'}^k, t_{e'})]_E$ be an edge in S'. By definition of the edges in Γ', the corresponding orbit of edge $e \in \Gamma$ is in $\pi(S) \subset \Gamma$. Thus there is an edge $[(a_{w'}^{c(e)}, t_e)]_E \in S \subset T$. In view of the definition of χ, the image of $\tilde{e}'$ by $\tilde{f}$ is $[(a_v^{nk}, a_v^{c(e)} t_e)]_E = [(a_v^{nk+c(e)}, t_e)]_E$. Since $a_v^{nk} \in \mathrm{Stab}(S)$ we have $\tilde{f}(\tilde{e}) = [a_v^{nk}(a_v^{c(e)}, t_e)]_E = a_v^{nk} \tilde{e}_S$ so it belongs to S.

The number of edges in S' is equal to the number of edges in S, and both can be computed with $[G_{\tilde{v}} : \mathrm{Stab}\, S]$ and the labels. Since $\tilde{f}$ is injective on the set of edges and sends S' to S, it induces a bijection between both sets so $\tilde{f}^{-1}(S) = S'$. This proves that T' is the tree $T^{\mathrm{Stab}\, S,S}$. □

Corollary 1.3.15. *Given T, H and S such as defined in lemma 1.1.22, the tree $T^{H,S}$ can be constructed algorithmically.* □

Lemma 1.3.16. *There is an algorithm which takes as input*

- *a marked graph of groups (Γ_T, ψ) with $\pi_1(\Gamma_T) \simeq G$, and universal cover T*
- *finite sets I_v^T associated to each vertex $v \in V(\Gamma_T)$ representing a family of allowed edge groups $\mathcal{A}$ such that $T \in \mathcal{D}^{\mathcal{A}}$*
- *a Whitehead graph $\mathrm{Wh}_T(\mathcal{G}, v)$ and an admissible cut of this graph for $\mathcal{A}$, where $\mathcal{G}$ is a finite collection of loxodromic elements of G*

and gives

- *a marked graph of groups Γ_S and universal cover $S \in \mathcal{D}^{\mathcal{A}}$*
- *sets I_v^S associated to vertices of Γ_S representing the same collection of allowed subgroups $\mathcal{A}$*

such that there exists a non-injective map $S \to T$, sending edge to edge or vertex, such that for every $g \in \mathcal{G}$, $\|g\|_S = \|g\|_T$. Moreover the map $S \to T$ can be chosen to be either a fold or a collapse.

Proof. The proof of lemma 1.2.17 gives a construction of such a tree S (or equivalently the associated graph of groups) by performing expansions and collapses on T. Lemmas 1.3.13 and 1.3.14 compute a new marked graph of groups for a collapse or an expansion, along with an isomorphism between fundamental groups which is compatible with the markings.

Finally we need to check that the sets which describe $\mathcal{A}$ can be computed, i.e. for any vertex group in Γ_S we need to find its maximal subgroups which belong to $\mathcal{A}$. For any vertex group H of Γ_S the change of markings given by lemmas 1.3.13 and 1.3.14 enables one to compute the image of H as a vertex subgroup in Γ_T, so using the family $(I_v^T)_{v\in\Gamma_T}$ one can compute the family $(I_v^S)_{v\in\Gamma_S}$. □

1.3.3 Description and termination of the algorithm

In this subsection we prove Theorem 1.3.3. Using the criterion given by theorem 1.2.14 we give an algorithm deciding, given G and $\mathcal{G} \subset G$ a finite collection of loxodromic elements, whether $\mathcal{G}$ is simple, and if so, determining a system of special factors $\mathcal{H}$ such that $\mathcal{G} \preceq \mathcal{H}$.

Note the following fact:

Lemma 1.3.17. *If $\mathrm{Wh}_T(\mathcal{G}, v)$ has an isolated vertex, then $\mathcal{G}$ is simple with respect to $\mathcal{D}$.*

Proof. In the quotient T/G the axis of g avoids some edge which defines a subforest of T/G containing the image of the axis of every $g \in \mathcal{G}$. □

Here is the description of the algorithm. We start with the tree $T_0 := T_G \in \mathcal{D}$ corresponding to the graph of groups Γ_G defining G. We may check immediately whether G is a solvable Baumslag-Solitar group, in which case no proper special factor exists. We suppose it is not the case. Start with $n = 0$:

1. Compute the axis of every $g \in \mathcal{G}$ in T_n. If some edge orbit does not intersect any axis, then its complementary subgraph is a subforest of Γ. It may have some components with elliptic fundamental group. Such components do not contain any element of $\mathcal{G}$. The set of components of the subforest with non-elliptic fundamental group gives a system of proper special factors $\mathcal{H}$ such that $\mathcal{G} \preceq \mathcal{H}$. The algorithm returns YES.

2. Compute Whitehead graphs for all vertices of T_n/G and check whether at least one of them has an admissible cut, using lemma 1.3.10. If none is found, then $\mathcal{G}$ is not simple and the algorithm returns NO.
3. If we find $v \in T_n/G$ such that $\mathrm{Wh}_{T_n}(\mathcal{G}, v)$ has an admissible cut, we compute a tree T_{n+1} such that there is a G-equivariant $f : T_{n+1} \to T_n$ sending edge to edge or to vertex with $\|g\|_{T_{n+1}} = \|g\|_{T_n}$ for every $g \in \mathcal{G}$. Lemma 1.3.16 ensures that this can be done. Start again step 1 with T_{n+1}.

Lemma 1.3.18. *The algorithm described above terminates. When it does, either it finds T^* such that $\mathrm{Axis}_{T^*}(g)$ does not cross some orbit of edges for every $g \in \mathcal{G}$, or it finds a proof that $\mathcal{G}$ is not simple.*

Proof. At each iteration of the second step we replace T_n by a tree T_{n+1} obtained by an expansion or an unfolding. Expansions and unfoldings of type A or B increase the number of orbits of edges by one while unfoldings of type C do not change this number. Suppose by contradiction that the algorithm does not terminate, yielding an infinite sequence $(T_n)_{n\in\mathbb{N}}$. Lemma 1.1.25 implies that the number of edges in T_n/G tends to infinity. Thus there exists $N \in \mathbb{N}$ such that the number of edges of T_N/G is strictly greater than $\sum_{g\in\mathcal{G}} \|g\|_{T_N} = \sum_{g\in\mathcal{G}} \|g\|_{T_0}$. Then there must be at least one orbit of edge avoided by the axis of every $g \in \mathcal{G}$, so the algorithm should have stopped at the N-th iteration: this is the contradiction we needed. □

1.4 Decreasing sequences of special factors

In this part, G is a non-elementary GBS group. Let $\mathcal{A}$ be a family of allowed edge subgroups of G. Let H be a special factor with respect to $\mathcal{D}^{\mathcal{A}}$. The induced deformation space $\mathcal{D}^{\mathcal{A}}_{|H}$ is the deformation space of H-trees of which elliptic groups and allowed edge groups are those of $\mathcal{D}$ which are contained in H.

In particular $\mathcal{D}_{|H}$ is the cyclic deformation space for H. If $\mathcal{A}_{|H} := \{A \cap H / A \in \mathcal{A}\}$, we have $\mathcal{D}^{\mathcal{A}}_{|H} = (\mathcal{D}_{|H})^{\mathcal{A}_{|H}}$.

Lemma 1.4.1. *A subgroup $K \subset H$ is a special factor of G with respect to $\mathcal{D}^{\mathcal{A}}$ if and only if K is a special factor of H with respect to $\mathcal{D}^{\mathcal{A}}_{|H}$.*

Proof. If T is a tree in $\bar{\mathcal{D}}^{\mathcal{A}}$ then its minimal H-invariant subtree T_H is a tree in $\bar{\mathcal{D}}^{\mathcal{A}}_{|H}$. Conversely, any H-tree in $\mathcal{D}_{|H}$ is a subtree of some G-tree in $\mathcal{D}$.

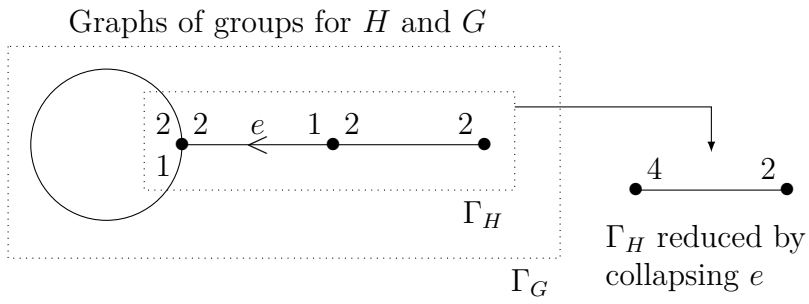

Figure 1.14 – Example: H is a special factor of G for $\mathcal{D}_{\text{red}}$ because G_e fixes the edge of the loop after reduction. On the contrary $(\mathcal{D}_{\text{red}})_{|H}$ is not the reduced space for H because G_e does not fix any edge in reduced H-trees (for example the universal cover of the graph of groups on the right).

Suppose $K \subset H$ is a special factor of G, then it is a vertex stabilizer of some vertex v in some $T \in \bar{\mathcal{D}}^{\mathcal{A}}$. The vertex v must belong to the minimal H-invariant subtree $T_H \in \bar{\mathcal{D}}^{\mathcal{A}}_{|H}$ since K is not an allowed edge group. Therefore K is a special factor for H with respect to $\mathcal{D}^{\mathcal{A}}_{|H}$.

Now suppose K is a special factor of H with respect to $\mathcal{D}^{\mathcal{A}}_{|H}$: there is a collapse $T \to \bar{T}$ with $T \in \mathcal{D}^{\mathcal{A}}_{|H}$ and such that K is a vertex stabilizer in $\bar{T}$. There is a tree $S \in \mathcal{D}^{\mathcal{A}}$ such that T is a subtree of S and there is a collapse $S \to \bar{S}$ such that the restriction to T is the collapse $T \to \bar{T}$ and the collapse is an isometry on $S \setminus T$. Edge stabilizers belong to $\mathcal{A}$. Thus K is a vertex stabilizer in $\bar{S}$ so it is a special factor for G. □

Remark 1.4.2. If $\mathcal{A} = \mathcal{A}_{\min}$ the family of subgroups which are bi-elliptic in reduced trees of $\mathcal{D}$, we have $\mathcal{D}^{\mathcal{A}} = \mathcal{D}_{\text{red}}$. However, the elements of $\mathcal{D}^{\mathcal{A}_{\min}}_{|H}$ are not necessarily reduced as H-trees of $\mathcal{D}_{|H}$. In fact some bi-elliptic groups appearing in some reduced G-trees might not be bi-elliptic in reduced H-trees (see figure 1.14 for an example). This means that $\mathcal{A}_{\min|H}$ *is not* the family of subgroups of H which are bi-elliptic in reduced H-trees of $\mathcal{D}_{|H}$.

Given a family $\mathcal{G}$ of loxodromic elements of G, the algorithm of Theorem 1.3.3 that we described can be iterated in order to find a $\preceq$-decreasing sequence $(\mathcal{H}_n)$ of systems of special factors with respect to $\mathcal{D}^{\mathcal{A}}$ such that for all $n \in \mathbb{N}$, $\mathcal{G} \preceq \mathcal{H}_n$. One can ask whether there exists a minimal system of special factors $\mathcal{H}_{\min}$ such that $\mathcal{G}$ is $\mathcal{H}_{\min}$-peripheral, and if the iteration of the algorithm eventually finds such a system. The answer is yes. We introduce a complexity on special factors which enables us to prove the existence of a minimal factor and that the algorithm stops.

Remark 1.4.3. Additional operations used to iterate the Whitehead algorithm are themselves algorithmic. We need two things. The first one is to be able to take a (possibly non-connected) subgraph $\Gamma' \subset \Gamma$, compute the corresponding subgroups $G'_1, \dots, G'_k$ and the minimal subtrees of these subgroup. Indeed the connected components of Γ' may have valence 1 vertices, in which case they do not represent the minimal subtree for the corresponding G'_i. The second one is to deduce the replacements for $\mathcal{D}^{\mathcal{A}}$ in every factor of the system. It consists in keeping all allowed edge groups which are contained in the special factor G'_i for every $i \in \{1, \dots, k\}$. In practice, for a factor G'_i corresponding to a connected component Γ'_i, this means we keep every I_v for which v belongs to Γ'_i.

For any non-elementary GBS group H we denote by $b_1(H)$ the first Betti number of any graph of cyclic groups with fundamental group H. This is an invariant of trees in the cyclic deformation space for H. We denote by $M(H)$ the set of conjugacy classes of big vertex stabilizers with respect to $\mathcal{A}_{\min}$. Define $m(H) := \#M(H)$. Vertex stabilizers are the same in all trees in $\mathcal{D}_{|H}$ so $m(H)$ is also well-defined.

We also introduce the following integer:

$$\sigma(H) = \sum_{K \in M(H)} i(K)$$

where $i(K)$ is an integer which we define as follows and which is linked to the peripheral structure of K (see [GL07, Definition 4.10])

Define $i(K) = [K : K']$ where $K' = \langle g \in K | g \text{ bi-elliptic in some } T \in \mathcal{D}^{\mathcal{A}_{\min}} \rangle$. This definition does not depend on any graph of groups for H. However one can compute $i(K)$ easily using the labels of a reduced graph of groups for H. Given such a graph, there exists exactly one vertex v corresponding to the conjugacy class K. Absolute values of labels at v are never 1 because K is big. Then $i(K)$ is the GCD of all labels at v.

Both $b_1(H), m(H), \sigma(H)$ only depend on the conjugacy class of H. For any special factor $H < G$ let $[H]$ be the conjugacy class of H in G.

Definition 1.4.4. We define the following complexity, which is a triple of non-negative integers, for any non-elementary GBS group H:

$$\mathcal{C}(H) = (b_1(H), m(H), \sigma(H))$$

It does not depend on the reduced graph of groups chosen to compute it nor on $\mathcal{A}$. We order complexities with lexicographic order.

For elementary GBS groups, we define the complexity to be $(0, 0, 0)$.

Proposition 1.4.5. *Let G be a GBS group and H a special factor of G. Then $\mathcal{C}(H) < \mathcal{C}(G)$.*

Remark 1.4.6. The proposition is true for any choice of $\mathcal{A}$. Since a special factor with respect to $\mathcal{D}^{\mathcal{A}}$ is a special factor with respect to $\mathcal{D}^{\mathbb{Z}}$, it suffices to prove Proposition 1.4.5 for $\mathcal{D}^{\mathbb{Z}}$.

Definition 1.4.7. Let Γ be a graph of groups and Γ' be a subgraph of Γ. We say that Γ is *reduced with respect to* Γ' if no collapse of Γ in the same deformation space has a subgraph with fundamental group $\pi_1(\Gamma')$.

This property is a minimality condition: if $\Gamma' \subset \Gamma$, up to collapsing some edges, we can obtain $\bar{\Gamma}' \subset \bar{\Gamma}$, where $\bar{\Gamma}$ is reduced with respect to Γ' and $\pi_1(\bar{\Gamma}') = \pi_1(\Gamma')$.

Remark 1.4.8. The definition implies that if an edge $e \in \Gamma$ has label $\lambda(e) = \pm 1$ and is not a loop, then $e \subset \Gamma \setminus \Gamma'$ and $o(e) \in \Gamma'$.

Proof of proposition 1.4.5. There exists a graph of groups Γ for G such that some subgraph $\Gamma_H \subset \Gamma$ has fundamental group H.

We may suppose that Γ is reduced with respect to Γ_H. In particular, Γ_H is reduced. We also suppose H is not an elementary GBS subgroup.

If at least one of the edges in $\Gamma \setminus \Gamma_H$ is non-separating then the first Betti number of Γ_H is strictly smaller than $b_1(\Gamma)$ so $\mathcal{C}(H) < \mathcal{C}(G)$.

If all edges in $\Gamma \setminus \Gamma_H$ are separating, then each connected component of $\Gamma \setminus \Gamma_H$ is a tree attached to Γ_H by a single vertex. Figure 1.15 illustrates this case.

In that case we first prove that the number of big vertex stabilizers (with respect to $\mathcal{A}_{\min}$) cannot increase.

Suppose $v \in \Gamma_H$ is a vertex whose stabilizer in H is big. If no label in $\Gamma \setminus \Gamma_H$ at v is ± 1, then $G_v < G$ is a big vertex stabilizer in Γ.

Otherwise, let e be an edge with $o(e) = v$ such that $\lambda(e) = \pm 1$. The edge e is not in Γ_H since v is big in Γ_H, thus e is separating. Furthermore no label at $t(e)$ is ± 1, so $G_{t(e)}$ is a big vertex group containing G_v. Thus every big stabilizer in H is a subgroup of a big stabilizer in G.

Moreover the subgroup $G_{t(e)} < G$ can contain at most one big vertex stabilizer of H. This implies $m(H) \leq m(G)$.

At least one vertex with two edges:

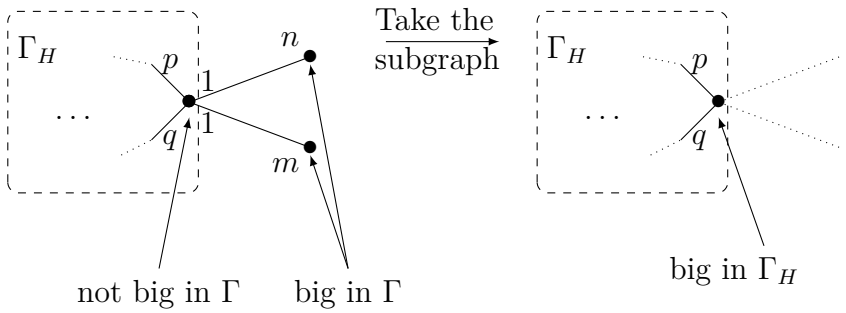

Only one edge on every vertex:

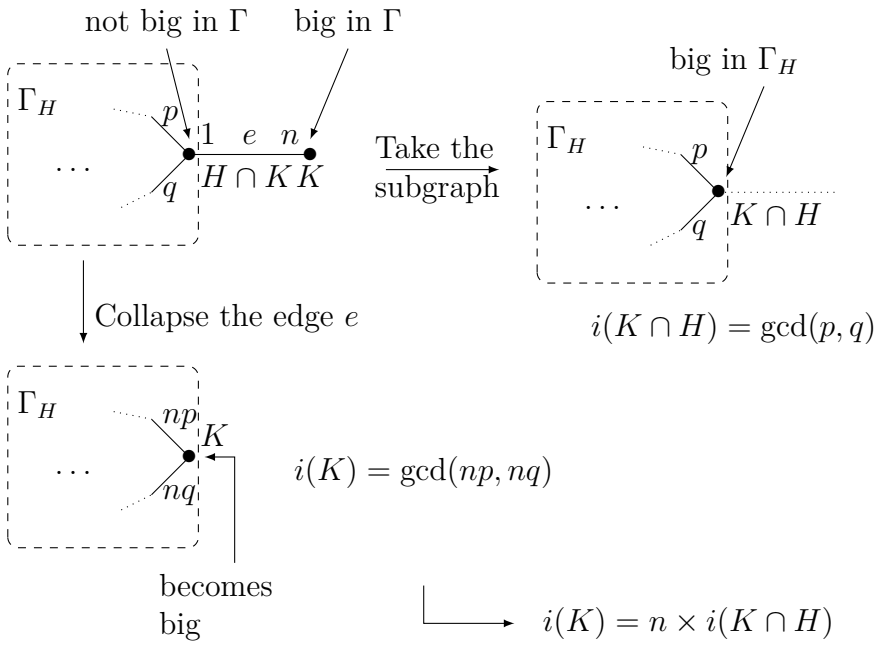

Figure 1.15 – Cases when all edges in $\Gamma \setminus \Gamma_H$ are separating.

Suppose at least one of the edges in $\Gamma \setminus \Gamma_H$ has both its label distinct from ± 1. One of its vertices v does not belong to Γ'. The stabilizer of v is big (see remark 1.1.20): no edge at v is a loop, and all labels at v are different from ± 1 because Γ is reduced with respect to Γ'. Thus $m(H) < m(G)$ so $\mathcal{C}(H) < \mathcal{C}(G)$ again.

Assume that all edges in $\Gamma \setminus \Gamma_H$ are separating and have one label equal to ± 1. In that case, all edges in $\Gamma \setminus \Gamma_H$ have one vertex in Γ_H, which carries the ± 1 label, and a valence 1 vertex. If at least two such edges are attached to the same vertex in Γ_H then Γ_H has strictly fewer big vertex stabilizer classes than Γ (see figure 1.15). Therefore we have to deal with the case where at most one edge of $\Gamma \setminus \Gamma_H$ is attached to each vertex of Γ_H.

In that case, compute σ in Γ and Γ_H. Only the latter is reduced. To reduce Γ we collapse all edges in $\Gamma \setminus \Gamma_H$. Thus we get a reduced graph Γ' which is similar to Γ_H but at some vertices, all labels are multiplied by a factor. At such vertices the GCD of all labels is also multiplied by the factor, so the corresponding $i_\Gamma(K)$ is greater than $i_{\Gamma_H}(K \cap H)$. Therefore $\sigma(H) < \sigma(G)$ so $\mathcal{C}(H) < \mathcal{C}(G)$. □

Remark 1.4.9. If H is a special factor of G with respect to $\mathcal{D}_{\mathrm{red}}(G)$, then $(b_1(H), m(H)) < (b_1(G), m(G))$. However if F is another special factor of G for $\mathcal{D}_{\mathrm{red}}(G)$, such that $F < H$, we do not necessarily have $(b_1(F), m(F)) < (b_1(H), m(H))$. In fact F may not be a special factor of H for $\mathcal{D}_{\mathrm{red}}(H)$ (see remark 1.4.2). Therefore this simplified complexity is not helpful to study a decreasing sequence of special factors, even with respect to $\mathcal{D}^{\mathcal{A}_{\min}}$.

We deduce:

Corollary 1.4.10. *Every decreasing sequence of special factors with respect to $\mathcal{D}^{\mathcal{A}}$ is stationary.*

Proof. Let $G \supset G_1 \supset \cdots \supset G_n \supset \ldots$ be a decreasing sequence of special factors of G. Either all the G_i are all non-elementary groups so Proposition 1.4.5 ensures that the sequence is stationary, or for some $n \in \mathbb{N}$ the group G_n is elementary. The elementary GBS groups do not have any non proper special factors so in that case the sequence is also stationary. □

Corollary 1.4.11. *Every $\preceq$-decreasing sequence of systems of special factors with respect to $\mathcal{D}^{\mathcal{A}}$ is stationary.*

Proof. This is a consequence of Corollary 1.4.10 and Kőnig's Lemma. □

Lemma 1.4.12. *If A, B are distinct special factors of G with respect to $\mathcal{D}^{\mathcal{A}}$, then either $A \cap B$ is elliptic or it is also a special factor of G with respect to $\mathcal{D}^{\mathcal{A}}$.*

Proof. Let A, B be as in the lemma. We will construct a tree in $\bar{\mathcal{D}}^{\mathcal{A}}$ in which $A \cap B$ is a vertex stabilizer.

There exists a G-tree $T \in \bar{\mathcal{D}}^{\mathcal{A}}$ where B is a vertex stabilizer. The group A acts on T; let T_A be the minimal subtree of T for this action. Subgroups of the form $hBh^{-1} \cap A$ fix a vertex in T_A. Consider the graph of groups T_A/A.

Let $S \in \bar{\mathcal{D}}^{\mathcal{A}}$ be such that A is the stabilizer of a vertex $v \in S$. We may suppose that all vertices with non cyclic stabilizers are in the orbit of v. We perform an expansion on S by replacing the vertices in the orbit of v by copies of T_A, which is possible since edge stabilizers of S are elliptic in T_A. We obtain a new tree R which has a vertex with stabilizer $B \cap A$.

We need to show that it can be obtained by collapse of a tree of $\mathcal{D}^{\mathcal{A}}$. We can perform an expansion of the vertex orbit of S by replacing it by $\hat{T}_A$, the minimal subtree of A in $\hat{T} \in \mathcal{D}^{\mathcal{A}}$ where $\hat{T} \to T$ is a collapse. All edge and vertex stabilizers of $\hat{T}_A$ are allowed in $\mathcal{D}^{\mathcal{A}}$. Then we get a tree $\hat{R} \in \mathcal{D}^{\mathcal{A}}$ which yields R by collapse. This shows that $A \cap B$ is either elliptic or a special factor of G with respect to $\mathcal{D}^{\mathcal{A}}$. □

Corollary 1.4.13. *Let $\mathcal{H}, \mathcal{H}'$ be two systems of proper special factors. Define $\mathcal{H} \wedge \mathcal{H}' := \{[H \cap H']/H \cap H' \text{ non elliptic }, [H] \in \mathcal{H}, [H'] \in \mathcal{H}'\}$. Then $\mathcal{H} \wedge \mathcal{H}'$ is a system of proper special factors.*

Proof. The elements of $\mathcal{H} \wedge \mathcal{H}'$ are all special factors according to Lemma 1.4.12. We need to check that they are all simultaneously vertex stabilizers in some tree $S \in \bar{\mathcal{D}}^{\mathcal{A}}$. The proof works like the proof of Lemma 1.4.12. Let $T \in \bar{\mathcal{D}}^{\mathcal{A}}$ (resp. T') be a tree in which every factor of $\mathcal{H}$ (resp. $\mathcal{H}'$) is a vertex stabilizer. We find minimal trees $T_i \subset T$ for H_i' for every $[H_i'] \in \mathcal{H}'$, then we blow up T' by replacing the vertex fixed by $gH_i'g^{-1}$ by gT_i for every $g \in G$. The result is a tree S in which the conjugacy classes of non-cyclic vertex stabilizers are the set $\mathcal{H} \wedge \mathcal{H}'$, and $S \in \bar{\mathcal{D}}^{\mathcal{A}}$. □

Corollary 1.4.14. *Let $\mathcal{G}$ be a finite collection of loxodromic elements of G. The set of systems of special factors $\mathcal{H}$ with respect to $\mathcal{D}^{\mathcal{A}}$ such that $\mathcal{G} \preceq \mathcal{H}$ admits a smallest element for $\preceq$.*

Proof. Note that this set is never empty since $\{G\}$ is itself a system of special factors with respect to $\mathcal{D}^{\mathcal{A}}$.

Corollary 1.4.11 ensures that any $\preceq$-decreasing sequence of systems of special factors is stationary. Thus there exists a system of special factors $\mathcal{H}$ such that $\mathcal{G} \preceq \mathcal{H}$ which is minimal for this property.

Let us show that it is unique. Let $\mathcal{H}'$ be another minimal special factor such that $\mathcal{G} \preceq \mathcal{H}'$. Let $g \in \mathcal{G}$: there is $H \in \mathcal{H}$ and $H' \in \mathcal{H}'$ such that $g \in H \cap H'$ so $\mathcal{H} \wedge \mathcal{H}'$ is not empty. By Lemma 1.4.12 it is a system of special factors of G.

By minimality of $\mathcal{H}$ we get $\mathcal{H} \subset \mathcal{H}'$, and conversely by minimality of $\mathcal{H}'$. By Remark 1.2.7 2. the relation $\preceq$ is an order so both systems are equal. □

Using Whitehead algorithm and corollary 1.4.14 we deduce that there exists an algorithm which finds the smallest system of special factors $\mathcal{H}$ such that $\mathcal{G} \subset G$ is $\mathcal{H}$-peripheral with respect to $\mathcal{D}^{\mathcal{A}}$:

Theorem 1.4.15. *There is an algorithm which takes as input*

— *a marked graph of groups Γ_G representing a group G*

— *a finite collection of loxodromic elements $\mathcal{G} \subset G$ as loops in Γ_G*

— *finite sets $(I_v)_{v \in V(\Gamma_G)}$ representing a collection of allowed edge groups $\mathcal{A}$*

and outputs the smallest system of special factors $\mathcal{H}$ of G with respect to $\mathcal{D}^{\mathcal{A}}$ such that $\mathcal{G}$ is $\mathcal{H}$-peripheral, as a finite collection of marked graph of groups.

Proof. We give the algorithm in the case where $\mathcal{G}$ has a single element. The algorithm consists in constructing a decreasing sequence of special factors containing g. In the case where $\mathcal{G}$ consists of more than one element, the algorithm would construct a decreasing sequence of systems of special factors.

Define $\mathcal{H}_0 := \{G\}$, $\Gamma_0 := \Gamma_G$ and $\mathcal{D}_0 := \mathcal{D}^{\mathcal{A}}$. We will construct a decreasing sequence of special factors $H_0 \supset \cdots \supset H_N$, where $g \in H_n$ for every $n \in \{0, \dots, N\}$. The space $\mathcal{D}_n$ is defined as the induced deformation space for H_n. It coincides with the induced deformation space for H_n seen as a special factor of some H_m for $m < n$. For every n, we will construct sets $(I_v^n)_{v \in V(\Gamma_n)}$ which represent the family of allowed edge groups of $\mathcal{D}_n$.

Here is the algorithm. Start with $i = 0$. Use Theorem 1.3.3 applied to Γ_i to decide whether g is contained in a special factor of H_i or not. If yes, then H_i is the minimal special factor containing g and the algorithm stops. Else the algorithm gives

— a special factor $H_{i+1} \subsetneq H_i$ such that $g \in H_{i+1}$

— a new graph of groups Γ_i' with $\pi_1(\Gamma_i') \simeq H_i$

— a subgraph of groups $\Gamma_{i+1} \subset \Gamma_i'$ such that $\pi_1(\Gamma_{i+1}) \simeq H_{i+1}$.

Then start again with $i + 1$ instead of i.

Applying this construction, we get a decreasing sequence of special factors with respect to $\mathcal{D}^{\mathcal{A}}$.

Corollary 1.4.10 guarantees that this sequence is stationary, which means that the algorithm stops eventually.The last factor obtained, given as a marked graph of groups, is the minimal special factor containing g. □

Chapitre 2

Detection of fully irreducible automorphisms in generalized Baumslag-Solitar groups

Introduction

Baumslag-Solitar groups are defined by

$$\mathrm{BS}(p,q) = \langle a, t | ta^p t^{-1} = a^q \rangle$$

They were introduced by Baumslag and Solitar in [BS62] as examples of non-Hopfian groups.

Generalized Baumslag-Solitar (GBS) groups are defined as fundamental groups of finite graphs of infinite cyclic groups (see Figure 2.1 for examples). Equivalently they are groups which act on a simplicial tree with infinite cyclic edge and vertex stabilizers and finite quotient.

In this paper we are interested in the outer automorphism group $\mathrm{Out}(G)$ for a GBS group G.

GBS groups naturally act on trees with infinite cyclic edge and vertex stabilizers. In general a GBS group admits infinitely many such actions, and there are infinitely many corresponding graphs of groups. In [For06] Forester introduced *deformation spaces* associated to a GBS group G. It consists in the space of G-trees which have the same elliptic groups as a given tree T. Here we consider the *cyclic deformation space* which is the space of all minimal actions of G on simplicial trees with infinite cyclic edge and vertex stabilizers, up to G-equivariant isomorphism and homothety (unless G is isomorphic to $\mathbb{Z}$, $\mathbb{Z}^2$ or the fundamental group of a Klein bottle). The cyclic deformation space $\mathcal{D}$ is the

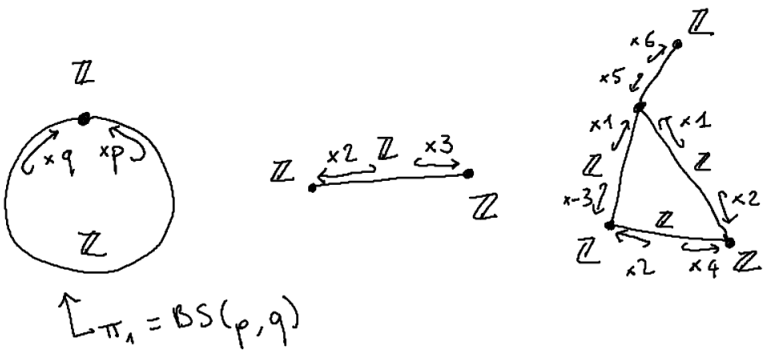

Figure 2.1 – Examples of graphs of groups with cyclic vertex and edge groups

GBS analogue of the outer space of Culler and Vogtmann CV_N for the free group F_N: outer space is a contractible space with a proper action of $\mathrm{Out}(F_N)$.

For a GBS group G, there is a natural action of $\mathrm{Aut}(G)$ on $\mathcal{D}$ by pre-composition of the action of G: if $\phi \in$
$Aut(G)$ and $(T, \rho) \in \mathcal{D}$ where $\rho : G \to \mathrm{Isom}(T)$ is the action, then $(T, \rho) \cdot \phi = (T, \rho \circ \phi)$. Since inner automorphisms act trivially this is actually an action of $\mathrm{Out}(G)$. In the rest of the paper, we distinguish actual automorphisms in $\mathrm{Aut}(G)$ and their outer class in $\mathrm{Out}(G)$: we denote the former with lower case letters and the latter with upper case letters.

Studying actions of G on trees with cyclic stabilizers is equivalent to studying *marked graphs of groups* (Γ, σ) where Γ is a graph of cyclic groups and σ is an isomorphism $G \simeq \pi_1(\Gamma)$.

An automorphism $\Phi \in \mathrm{Out}(F_N)$ is called *fully irreducible* when no conjugacy class of free factors is Φ-periodic. In the case of GBS groups, free factors are not relevant but there is an analogue called *special factor*, which we develop in Chapter 1. A special factor is a subgroup of G which is the fundamental group of a subgraph of groups in some graph of cyclic groups for G. We can define a fully irreducible automorphism $\Phi \in \mathrm{Out}(G)$: it is an automorphism which has no periodic conjugacy class of special factor.

In [BH92] Bestvina and Handel prove that any fully irreducible element $\Phi \in \mathrm{Out}(F_N)$ admits a *train track representative*, which is an action on a tree $T \in \mathcal{D}$ with a ϕ-equivariant map $f : T \to T$ for some $\phi \in \Phi$ whose iterates stretch equally all edges and send them on geodesic paths. The set of train track trees behave like a translation axis for Φ in CV_N. Train track representatives are a key tool for studying fully irreducible automorphisms and they can be computed. However it is not known whether train track representatives

always exist for fully irreducible automorphisms of a GBS group in general, although there are many examples. When G has no *non-trivial integer modulus* (see for example [Lev07] for a definition of modulus) then there is a bound on the number of edges in graphs of cyclic groups for G ([For06]). The results in [Mei15] imply that fully irreducible automorphisms of G in this case admit train track representatives. Yet we do not know whether such groups always admit interesting fully irreducible automorphisms.

Another example is $G = \mathrm{BS}(p, pn)$ for $p > 1$. In [Bou16] Bouette proved that in that case, all automorphisms are reducible since there is a globally invariant conjugacy class of special factors. However by replacing $\mathcal{D}$ by another deformation space allowing greater vertex stabilizers, namely those in the invariant conjugacy class, there are fully irreducible automorphisms and they admit train tracks.

In this paper we adapt an algorithm given by Kapovich in [Kap14] and [Kap19] which, given an automorphism Φ which is *atoroidal*, i.e. has no periodic conjugacy class, and a train track representative for Φ, decides whether or not Φ is fully irreducible. The atoroidal condition can actually be weakened. In the original papers by Kapovich the algorithm applies to any automorphism since there is an algorithm in [BH92] which either finds a train track representative for Φ or gives a proof that Φ is not reducible.

A non-solvable GBS group is a GBS group which is neither isomorphic to $\mathbb{Z}$, $\mathbb{Z}^2$, the fundamental group of a Klein bottle nor $\mathrm{BS}(1, n)$ for some $n \in \mathbb{Z} \setminus \{0\}$.

A conjugacy class $[g] \in G$ is *loxodromic* if there exists $T \in \mathcal{D}$ (equivalently for any $T \in \mathcal{D}$) the action of g on T is loxodromic.

A loxodromic conjugacy class $[g]$ with $g \in G$ is *pseudo-periodic* for $\Phi \in \mathrm{Out}(G)$ if $\|\phi^n(g)\|$ is bounded for $n \in \mathbb{N}$, for some automorphism $\phi \in \Phi$. While in a free group context this would ensure that the conjugacy class of g is actually periodic, it is not the case here. An outer class $\Phi \in \mathrm{Out}(G)$ is *pseudo-atoroidal* if it has no pseudo-periodic hyperbolic conjugacy class.

The paper is centred around the two following theorems, which are independant.

Theorem A. *There is an algorithm which takes a non-solvable GBS group G, a pseudo-atoroidal automorphism $\Phi \in \mathrm{Out}(G)$ and a train track representative $f : T \to T$ for Φ, and decides whether Φ is fully irreducible.*

Theorem B. *There is an algorithm which takes a non-solvable GBS group G and an outer automorphism $\Phi \in \mathrm{Out}(G)$ with a train track representative $f : T \to T$ and decides whether Φ is pseudo-atoroidal.*

Structure of pseudo-periodic conjugacy classes and subgroups. The fixed group for an automorphism $\phi \in \mathrm{Aut}(G)$ is the subgroup of G whose elements are fixed for ϕ. When ϕ is an automorphism of a free group, the fixed group of ϕ is always finitely generated ([Ger87]). In GBS groups, fixed groups are not always finitely generated: see Example 2.0.1. The notion of pseudo-periodic subgroup associated to an automorphism $\phi \in \mathrm{Aut}(G)$ seems to be more interesting.

Examples 2.0.1. Let $G := \mathrm{BS}(2,4) = \langle a, t | ta^2t^{-1} = a^4 \rangle$. Let $\phi \in \mathrm{Aut}(G)$ be the inner automorphism defined by $g \mapsto a^2ga^{-2}$. The fixed group $\mathrm{Fix}(\phi)$ is the subgroup generated by $\{t^{-k}a^2t^k, k \in \mathbb{Z}\}$ and it is not finitely generated. In fact, let $T \in \mathcal{D}$ be the standard tree for $\mathrm{BS}(2,4)$ of Figure 2.1. Let v be the unique vertex of T with stabilizer $\langle a \rangle$. Let e be the edge from v to tv. The minimal invariant subtree for $\mathrm{Fix}(\phi)$ is the subset of T which can be reached from v by an edge path $\gamma := e_1, \ldots, e_n$ such that for every $i \in \{1, \ldots, n\}$ the subpath $e_1, \ldots, e_i$ does not contain more edges in $G \cdot e$ than in $G \cdot \bar{e}$.

The quotient of the minimal subtree by $\mathrm{Fix}(\phi)$ has infinite diameter. Therefore $\mathrm{Fix}(\phi)$ is not finitely generated.

In free groups, pseudo-periodic conjugacy classes are actually periodic. However, in a GBS group, this is not automatically true. In Chapter 1 we proved the existence of a minimal special factor containing g for every loxodromic $g \in G$. A conjugacy class is *simple* if its minimal special factor is not G. If $[g]$ is pseudo-periodic, then the conjugacy class of the minimal factor containing g is periodic: if g is simple, then G has a proper periodic conjugacy class of special factors. Thus we have:

Lemma C. *If there exists a simple pseudo-periodic conjugacy class for $\Phi \in \mathrm{Out}(G)$, then Φ is not fully irreducible.*

The strategy for finding pseudo-periodic conjugacy classes is based on *Nielsen paths*, just as the study of periodic conjugacy classes in free groups. Let $f : T \to T$ be a train track representative for $\Phi \in \mathrm{Out}(G)$. A *Nielsen path* is a finite path $[x, y]$ in T such that there exists $g \in G$ such that $f(x) = gx$ and $f(y) = gy$. More generally a periodic Nielsen path is a Nielsen path for f^n for some $n \geq 1$. They can be split into concatenations of *periodic indivisible Nielsen paths* (pINP). The link between pseudo-periodic conjugacy classes for Φ and periodic indivisible Nielsen paths for f is given by:

Proposition D. *Let $g \in G$ be an element whose conjugacy class is pseudo-periodic for $\Phi \in \mathrm{Out}(G)$. Suppose $f : T \to T$ is a train-track representative for Φ. Then the axis of g in T is a concatenation of periodic indivisible Nielsen paths.*

Consider a periodic indivisible Nielsen path $\eta \subset T$. Define $VY(\eta)$ as the set of points of T which can be joined to one endpoint of η by a concatenation of periodic indivisible Nielsen paths. By Proposition D the axis of $g \in G$ such that the conjugacy class $[g]$ is pseudo-periodic is contained in the convex hull of $VY(\eta)$ for some periodic indivisible Nielsen path η and that that $g \in \mathrm{Stab}(VY(\eta))$.

Like in the case of free groups, there are finitely many G-orbits of periodic indivisible Nielsen paths, and they can be computed:

Theorem E. *Let $\Phi \in \mathrm{Out}(G)$ and let $f : T \to T$ be a train track map for Φ. There are finitely many orbits of periodic indivisible Nielsen paths for f.*

Furthermore there is an algorithm which finds all orbits of pINPs.

Consequently there exist finitely many sets $VY(\eta)$ where η is a pINP up to translation, so there exist finitely many conjugacy classes of subgroups $\mathrm{Stab}(VY(\eta))$ which could contain elements whose conjugacy class is pseudo-periodic:

Theorem F. *Let $\Phi \in \mathrm{Out}(G)$. There exist finitely many subgroups*

$$G_1 = \mathrm{Stab}(VY(\eta_1)), \ldots, G_k = \mathrm{Stab}(VY(\eta_k))$$

in G, well defined up to conjugacy, for which one can compute a finite set of generators, such that:

- *If a conjugacy class $[g]$ with $g \in G$ is pseudo-periodic for Φ then there exists $i \in \{1, \ldots, k\}$ such that a conjugate of g belongs to G_i.*
- *Conversely, for all $i \in \{1, \ldots, k\}$, all loxodromic elements of G_i have a pseudo-periodic conjugacy class.*

Theorem B follows.

Given an actual automorphism $\phi \in \mathrm{Aut}(G)$ and a tree $T \in \mathcal{D}$, one can define the *pseudo-periodic subgroup* associated to G_ϕ as the subgroup

$$G_\phi := \{g \in G / d_T(*, \phi^n(g)*) \text{ is bounded for } n \in \mathbb{N}\}$$

where $*$ is any base point in T. The group G_ϕ does not depend on the choice of the base point $*$, nor on the choice of $T \in \mathcal{D}$.

Pseudo-periodic subgroups of GBS groups are related to fixed groups for automorphisms of free groups. There is also a link with Nielsen paths and pseudo-periodic conjugacy classes:

Theorem G. *Let $\Phi \in \mathrm{Out}(G)$ and let $f : T \to T$ be a train track representative for Φ. Let $\psi \in \mathrm{Aut}(G)$ such that $\psi \in \Phi^k$ for some $k \in \mathbb{N}$. If G_ψ contains a loxodromic element, there exists a pINP $\eta \subset T$ such that*

$$G_\psi \subset \mathrm{Stab}(VY(\eta))$$

Moreover, for any pINP η such that $\mathrm{Stab}(VY(\eta))$ contains a loxodromic element, there exists $k \in \mathbb{N}$ and $\psi \in \Phi^k$ such that $\mathrm{Stab}(VY(\eta)) = G_\psi$.

This indicates that the subgroups $\mathrm{Stab}(VY(\eta))$ for η a pINP are the maximal subgroups among the pseudo-periodic subgroups which contain loxodromic elements.

Periodic special factors and stable lamination. Given an outer automorphism $\Phi \in \mathrm{Out}(G)$ with a train track representative $f : T \to T$, one can define the *stable lamination* Λ^+ associated to T as the G-invariant collection of bi-infinite geodesics of T obtained by taking the closure of the set of translates of iterates $f^n(e)$ for an edge $e \in T$.

Theorem A is a consequence of the following criterion extending [Kap14]:

Theorem H. *Let G be a non-solvable generalized Baumslag-Solitar group. Let $\Phi \in \mathrm{Out}(G)$ be an automorphism of G with a primitive train track $f : T \to T$ with no simple pseudo-periodic element. Then Φ is fully irreducible if and only if all Whitehead graphs $\mathrm{Wh}_T(\Lambda^+, v)$ are connected.*

A train track map is *primitive* when the associated transition matrix is primitive.

We actually prove a slightly stronger version of this criterion, using restricted deformation spaces instead of the standard deformation space (see [GL07, Def. 3.12]): restricted deformation spaces have a restriction on allowed edge groups as well as on vertex groups.

From this criterion we deduce the algorithm of Theorem A. The algorithmicity here comes from the fact that $\Lambda^+(f)$ is quasi-periodic, which guarantees that the computation of the turns taken by the lamination takes only finitely many steps. The computation of Whitehead graphs is possible since the trees in $\mathcal{D}$ are locally finite.

Note that the criterion given by Theorem H does not require that Φ be pseudo-atoroidal: Theorem A actually applies to automorphisms with non-simple pseudo-periodic

conjugacy classes. However we do not know how to algorithmically determine whether an automorphism satisfies this hypothesis. We can decide whether there exist pseudo-periodic conjugacy classes, and for a given element we can decide whether it is algorithmic using the results of Chapter 1, but the set of pseudo-periodic conjugacy classes could be complicated.

In order to answer this question, we would need to know whether one of the subgroups $\mathrm{Stab}(VY(\eta))$ defined before contains a simple element. Without any algorithmic answer to this question, we need the pseudo-atoroidal condition to assemble Theorems B and A.

In Section 2.1 we define irreducible automorphisms and train track representatives, as well as the stable lamination which comes with a train track map. In Section 2.2 we prove Theorem H. Section 2.3 is about pseudo-periodic conjugacy classes and how they may be understood using Nielsen paths; pseudo-periodic subgroups and their computation is developed in Section 2.4. These two sections can be read independently of Section 2.2. Finally Section 2.5 puts together the arguments for Theorem A.

2.1 Special factors, automorphisms of GBS groups

A Generalized Baumslag-Solitar group (GBS) is a group which admits a minimal action on a simplicial tree with infinite cyclic vertex and edge stabilizers. Equivalently it is a group isomorphic to the fundamental group of a finite graph of infinite cyclic groups on both vertices and edges.

GBS groups isomorphic to $\mathbb{Z}$, $\mathbb{Z}^2$ and the fundamental group of a Klein bottle are called *elementary*. Along with Baumslag-Solitar groups $\mathrm{BS}(1,n)$ for $n \in \mathbb{Z} \setminus \{0\}$, they are the solvable GBS groups. In this paper we will always assume that G is a non-solvable GBS group.

A *G-tree* is a simplicial tree with an action of G by graph isomorphisms. It can be endowed with a metric, in which case we demand that the action be by isometries.

The set of vertices of a tree T will be denoted by $V(T)$ and the set of edges by $E(T)$. The initial vertex of an edge e is $o(e)$ and its terminal edge is $t(e)$. The opposite edge of e is $\bar{e}$.

We identify every tree T with its geometric realization. A *path* in a tree T is a continuous map from an interval to T. We will assume paths are linear on edges. We will frequently identify paths with their image in T. A *tight* path is a path which has no backtracking. An *edge path* is a path whose image can be described by whole edges $e_1, \dots, e_k$ such that for $1 \leq i < k$, $t(e_i) = o(e_{i+1})$.

Let T be a G-tree. A subgroup $H < G$ is *elliptic* in T if it fixes a point. It is *bi-elliptic* if it fixes two distinct points. When H is not elliptic and contains a loxodromic element, there exists a minimal H-invariant subtree which we denote by T_H.

A *deformation space* introduced by Forester in [For06] is the space of all minimal G-trees which have the same elliptic subgroups, up to G-equivariant isometry and homothety.

The *cyclic deformation space* $\mathcal{D}$ is the deformation space of all G-trees where elliptic subgroups are all isomorphic to $\mathbb{Z}$. This is well defined. Indeed there exists a G-tree T such that all edge and vertex stabilizers in T are infinite cyclic. Define $\mathcal{D}_T$ as the deformation space containing T. For non-elementary GBS groups the set of elliptic subgroups (i.e. subgroups which fix at least a point) is independent of the tree and has an algebraic characterization (see [Lev07] for example) so we set $\mathcal{D} := \mathcal{D}_T$.

Let $\mathcal{A}$ be a family of elliptic subgroups in $\mathcal{D}$, invariant by conjugacy and by passing to a subgroup. The *restricted deformation space* $\mathcal{D}^{\mathcal{A}}$ is the subspace of trees in $\mathcal{D}$ where all bi-elliptic groups belong to $\mathcal{A}$.

Let T be a G-tree. Let Y be a G-invariant proper subforest of T. Define the equivalence relation $\sim$ on T as the minimal G-equivariant equivalence relation such that all points of a connected component of Y are equivalent. The quotient map $\pi_Y : T \to T/\sim$ is the *collapse* of the subforest Y. The tree $T/\sim$ belongs to the same deformation space as T if and only if Y does not contain the axis of any loxodromic element of G.

A subgraph of the graph of groups is *collapsible* if its pre-image in the universal cover does not contain the axis of a loxodromic element. Equivalently, when collapsing the corresponding subforest, the stabilizers of the new vertices are elliptic subgroups of G.

Remark 2.1.1. Collapsibility can be checked using labels. To determine if a subgraph $\Gamma_0 \subset \Gamma$ is collapsible, first of all obtain a new graph from Γ_0 by deleting all edges with a valence 1 vertex carrying the label ± 1. This yields a graph of groups Γ_1. Then repeat with Γ_1. Since the number of edges decreases the procedure will eventually stop, yielding either a graph with no edge, or a graph where all valence 1 vertices have labels different from ± 1. The subgraph Γ_0 is collapsible if and only if the first case happens.

Let $\mathcal{D}^{\mathcal{A}}$ be a restricted deformation space, possibly $\mathcal{D}$ if all edge groups are allowed. A *special factor* with respect to $\mathcal{D}^{\mathcal{A}}$ is a subgroup H of G with the following property (see Chapter 1 for more details). There exists a collapse map $\pi : S \to \bar{S}$ such that $S \in \mathcal{D}^{\mathcal{A}}$ and H is a vertex stabilizer in $\bar{S}$.

A *marked graph of groups* is a pair (Γ, σ) where Γ is a graph of groups and $\sigma : G \to$

$\pi_1(\Gamma)$ is an isomorphism. Bass-Serre theory gives a correspondance between marked graphs of groups and G-trees ([Ser77]).

Marked graphs of groups offer a nice point of view regarding special factors. A subgroup $H < G$ is a special factor if and only if there exists a marked graph of groups (Γ, σ) and an identification $\sigma : G \to \pi_1(\Gamma)$ such that H identifies to the fundamental group of a subgraph of Γ. In this definition, Γ is the quotient S/G with $S \in \mathcal{D}^{\mathcal{A}}$ such that H is a vertex stabilizer in a collapse $\bar{S}$.

In the rest of the paper we fix $\mathcal{A}$, so we omit to specify it when referring to special factors.

Definition 2.1.2. The outer automorphism $\Phi \in \mathrm{Out}(G)$ is *reducible* with respect to $\mathcal{D}^{\mathcal{A}}$ if there exists a special factor $A \subset G$ with respect to $\mathcal{D}^{\mathcal{A}}$ such that the conjugacy class of A is invariant by Φ.

The outer automorphism Φ is called *fully irreducible* if for all $n \in \mathbb{N}$, Φ^n is not reducible.

Remark 2.1.3. Since the notion of special factors depends on $\mathcal{A}$, so does the notion of reducible automorphisms. Actually if $\mathcal{A} \subset \mathcal{A}'$ then special factors with respect to $\mathcal{A}$ are special factors with respect to $\mathcal{A}'$, so an automorphism which is reducible for $\mathcal{A}$ is also reducible for $\mathcal{A}'$. The strongest notion of fully irreducible automorphism is achieved with $\mathcal{D}^{\mathcal{A}} = \mathcal{D}$. We have not explored the possible differences caused by the choice of $\mathcal{A}$ yet.

2.1.1 Computations in GBS groups and trees

In the rest of the paper we will need to perform some algorithmic operations on GBS groups and their automorphisms. In this subsection we justify why and how we can do them. The main operations are summed up by Proposition 2.1.4.

In order to deal with the group G algorithmically, we consider elements of G as words in a graph of groups. The trees on which G act also admit a description as a set of words. The definitions in this section are standard, though they are given in the special case of GBS groups whose graphs of groups have cyclic edge and vertex groups, hence the simpler relations. A more detailed description is available in Chapter 1, Section 1.3.

Let Γ be a graph of groups with cyclic edge and vertex groups. Let $(a_v)_{v \in V(\Gamma)}$ be generators for every vertex group G_v of Γ and $(a_e)_{e \in E(\Gamma)}$ be a choice of generators for the edge groups. Each inclusion $G_e \hookrightarrow G_{t(e)}$ is given by a nonzero integer $\lambda(\bar{e})$ such that a_e is sent to $a_{t(e)}^{\lambda(\bar{e})}$. The integers $\lambda(e)$ is the *label* of e.

Let $(t_e)_{e\in E(\Gamma)}$. The *Bass group* $B(\Gamma)$ is the group generated by the elements a_v and t_e for $v \in V(\Gamma)$ and $e \in E(\Gamma)$, and with the following relations:

— $t_e = t_{\bar{e}}^{-1}$ for $e \in E(\Gamma)$

— for all $e \in \Gamma$, with labels $\lambda(e) = p$ and $\lambda(\bar{e}) = q$, we have $a_{o(e)}^p = t_e a_{t(e)}^q t_e^{-1}$

A *path* in the graph of groups is a word $a_{v_0}^{k_0} t_{e_1} a_{v_1}^{k_1} \dots t_{e_n} a_{v_n}^{k_n}$ where $v_{i-1} = o(e_i)$, $v_i = t(e_i)$ for $1 \leq i \leq n$ and $k_i \in \mathbb{Z}$ for $0 \leq i \leq n$. It is a *loop* if $v_0 = v_n$; in that case we say that it is *based* at v_0. Paths represent elements of $B(\Gamma)$. The integer n is the *length* of the path.

Fix a base point $v_0 \in V(\Gamma)$. The *fundamental group* of Γ at vertex v_0 is the subgroup $\pi_1(\Gamma, v_0) < B(\Gamma)$ whose elements are represented by loops in Γ based at v_0. The group is independant of the choice of v_0 up to isomorphism.

A path $w := a_{v_0}^{k_0} t_{e_1} a_{v_1}^{k_1} \dots t_{e_n} a_{v_n}^{k_n}$ in Γ is *reduced* if for all $1 \leq i \leq n-1$, $e_i = \bar{e}_{i+1}$ implies that $\lambda(\bar{e}_i)$ does not divide k_i.

Given a path w, one can algorithmically compute a reduced path which represents the same element of $B(\Gamma)$. The reduced path is not necessarily unique, however its length is unique.

The *universal cover* T_{v_0} of the graph of groups Γ is a tree which can be constructed as follows: the set of vertices is

$$\tilde{V} = \{\text{paths in } \Gamma \text{ with initial vertex } v_0\}/\sim$$

where $\gamma \sim \gamma'$ if γ and γ' have the same terminal vertex $v_i \in \Gamma$ and $\gamma^{-1}\gamma' \in G_{v_i}$ as an element of $B(\Gamma)$.

The oriented edges of the universal cover are defined as follows:

$$\tilde{E} = \left\{(\alpha, at_e)/\ \alpha \text{ path in } \Gamma \text{ from } v \text{ to } v', e \in \mathcal{E}_{v'}, a \in G_{v'}\ \right\}/\sim$$

where $\mathcal{E}_{v'}$ is the set of edges with origin v'. The equivalence relation $\sim$ is defined by $(\alpha, at_e) \sim (\alpha', a't_{e'})$ if and only if $e = e'$ and $a^{-1}\alpha^{-1}\alpha' a' \in i_{\bar{e}}(G_e)$. The opposite edge of (α, at_e) is the edge $(\alpha \cdot at_e, \bar{t}_e)$.

The group $\pi_1(\Gamma, v)$ acts on $\tilde{V}$ and $\tilde{E}$ by left concatenation of the paths. The quotient of T_{v_0} by this action is Γ.

Proposition 2.1.4. *Let G be a GBS group given as a finite set of generators and a marking $G \simeq \pi_1(\Gamma, v)$ where Γ is a graph of group with cyclic edge and vertex stabilizers and v is a vertex in Γ.*

One can algorithmically solve the following problems:

(i) word problem: given an element $g \in G$ given as a loop in the graph of groups, find out whether g is the identity

(ii) given $x \in \tilde{V}$ (resp. $e \in \tilde{E}$) and $g \in G$ given as a loop, compute the image gx

(iii) given vertices $x, x' \in \tilde{V}$ (resp. edges $e, e' \in \tilde{E}$) given as paths in the graph of groups, decide if $x \sim x'$ (resp. $e \sim e'$), i.e. if the paths x, x' represent the same point of T_v

(iv) given x, y in $\tilde{V}$, find out whether x, y belong to the same orbit, and when they do, find $g \in G$ such that $y = gx$

(v) given two pairs $(x, x'), (y, y')$ of elements of $\tilde{V}$, find out whether there exists $g \in G$ such that $(y, y') = (gx, gx')$

(vi) given $x \in \tilde{V}$ and $n \in \mathbb{N}$, list all edge paths with length n and first vertex x

(vii) given an edge $e \in E(\Gamma)$, compute the graph of groups Γ' obtained from Γ by subdividing e

Proof. (i) Using the identification of the set of generators of G with loops in Γ, one can compute a loop w in Γ for g. By [Ser77, 5.2, Theorem 11] a word w represents the trivial word if and only if it can be reduced to the trivial loop $1 = a_{v_0}^0$, or equivalently if every reduced form is the trivial loop. Since computing a reduced loop is algorithmic, the word problem can be solved.

(ii) To compute the image gx of $x \in V(T)$ given as a path in Γ based at v_0, it suffices to compute a loop for g based at v_0 and concatenate it on the left of the path representing x.

(iii) Let γ, γ' be paths in Γ representing x, x'. Their first vertex is v_0 and their last vertex is a vertex $v_i \in V(\Gamma)$. The concatenation $\gamma^{-1} \cdot \gamma'$ is a loop based at v_i. It represents an element of $G_{v_i} \subset B(\Gamma)$ if and only if there exists a reduced loop for it with length 0 (or equivalently, if every reduced loop for $\gamma^{-1}\gamma'$ has length 0).

(iv) The points x, y belong to the same orbit in T if and only if the paths representing them have the same terminal vertex in $V(\Gamma)$. The element $g \in G$ represented by the loop yx^{-1} satisfies $gx = y$.

(v) We can assume that there exists $g \in G$ such that $gx = y$ and find such a g. The set of elements of G such that $gx = y$ is $\{gu, u \in \mathrm{Stab}(x)\}$. Let γ_x be a path representing

x and let v be its terminal vertex in $V(\Gamma)$. By definition of $\tilde{V}$, $\mathrm{Stab}(x) = \{a_x^i, i \in \mathbb{Z}\}$ with $a_x := \gamma_x a_v \gamma_x^{-1}$. Since a_v is elliptic on T_{v_0} and the tree T_{v_0} is locally finite, the orbit $\langle a_v \rangle \cdot x'$ is finite so there exists $i \geq 1$ such that $a_x^i x' = x'$, and one can algorithmically find such an i by (iii). Thus the pairs (x, x') and (y, y') are in the same orbit if and only if there exists $0 \leq j < i$ such that $g a_x^j x' = y'$. This can be checked algorithmically.

(vi) This can be done by induction. For $n = 1$ it suffices to list all edges with origin x: if γ_x is a path in Γ representing x and v_x is the image of x in $V(\Gamma)$, then the edges are

$$\{(\gamma_x, a_{v_x}^i t_e, e \in E(\Gamma) \text{ s.t. } o(e) = v_x, 0 \leq i < \lambda(e)\}$$

For $n > 1$, list all edge paths with length $n - 1$. The edge paths with length n are the paths $\tilde{\gamma} \cdot \tilde{e}$ where $\tilde{\gamma}$ is an edge path with length $n - 1$ starting with x, $\tilde{e} \in \tilde{E}$ is an edge whose origin is the last vertex of $\tilde{\gamma}$, and the last edge of $\tilde{\gamma}$ and $\tilde{e}$ are not opposite.

(vii) The subdivision consists in replacing e by edges e', e'' with $o(e') = o(e)$, $t(e'') = t(e)$, and $t(e') = o(e'')$ is a new vertex. Labels are the following: $\lambda(e') = \lambda(e)$, $\lambda(\bar{e}') = \lambda(e'') = 1$, $\lambda(\bar{e}'') = \lambda(\bar{e})$. The new marking must also be computed; the operation consists in replacing t_e in paths in Γ by $t_{e'} 1 t_{e''}$. □

A group automorphism $\phi : G \to G$ can be described using a finite set of generators for G. We also assume that we know a description of the inverse of the automorphism $G \to \pi_1(\Gamma)$ in terms of the generators of $\pi_1(\Gamma)$ described above. Thus one can compute the image of any element of $\pi_1(\Gamma)$ by ϕ and the image is a loop in Γ.

Maps between trees can be described using the same tool. For a G-equivariant map between trees (or a ϕ-equivariant map; the G-equivariant case is the case $\phi = \mathrm{id}$) sending vertex to vertex, it suffices to give the image of every orbit of vertex and edge in the tree. Let $f : T \to T$ be a G-equivariant map and let $\mathcal{V} \subset V(T)$ be a set of representatives for $V(T/G)$, let $\mathcal{E} \subset E(T)$ be a set of representatives for $E(T/G)$:

— the image of a vertex $v \in \mathcal{V}$ is a path $f(v) \in \tilde{V}$

— the image of an edge $e \in \mathcal{E}$ is a path in the universal cover: it may be a single vertex or several edges.

For another vertex gv (resp. edge ge) with $g \in G$ described as a loop in $\pi_1(\Gamma)$, the image is $\phi(g)f(v)$ (resp. $\phi(g)f(e)$) where $\phi(g)$ and $f(v)$ are paths which can be computed separately.

2.1.2 Irreducible automorphisms and train track representatives

Let T, T' be metric G-trees. Let d be the distance on T and d' be the distance on T'. Let $f : T \to T'$ be a map sending vertex to vertex and edge to non-backtracking edge path. The *Lipschitz constant* of f is $\mathrm{Lip}(f) := \sup_{x \neq y \in T} \frac{d'(f(x), f(y))}{d(x,y)}$.

In the rest of the paper, we will always assume that maps between trees send vertex to vertex and send edges to non-backtracking edge paths. We need not assume that maps are linear on edges but for $e \in E(T)$ and $e' \in f(e)$ we will assume that it is linear on $f^{-1}_{|\{e\}}(\{e'\})$.

Let $v \in T$ be a vertex. A *turn* at v is an unordered pair of distinct edges $\{e, e'\}$ with origin v. If $e = e'$ we call it a *degenerate turn*. A non-backtracking path γ *crosses* a turn $\{e, e'\}$ if e and e' appear in γ.

Fix a map $f : T \to T$. A turn $\{e, e'\}$ is *illegal* if there exists $n \geq 1$ such that $f^n(e)$ and $f^n(e')$ are paths with a common prefix of nonzero length. Otherwise the turn is *legal*. A non-backtracking path γ is a *legal path* if every turn crossed by γ is legal.

Definition 2.1.5. We say that f is a *train track map* if f sends every edge $e \in E(T)$ to a legal path. It is a *metric train track map* if in addition the stretch factor on every edge is uniform and equal to $\mathrm{Lip}(f)$, i.e. $\mathrm{len}(f(e)) = \mathrm{Lip}(f)\,\mathrm{len}(e)$.

When f is a train track map, for every $e \in E(T)$ and every $n \in \mathbb{N}$ the path $f^n(e)$ is a geodesic.

The *bounded cancellation constant* $\mathrm{BCC}(f)$ is a constant introduced in [Coo87] in the case of free groups, with the following property. Let α, β be legal paths and let $\alpha \cdot \beta$ be their concatenation. The path $\alpha \cdot \beta$ may not be legal and subsegments of $f(\alpha)$ and $f(\beta)$ may be equal. However the length of the common subsegments is bounded by $\mathrm{BCC}(f)$. The constant exists for the following reason. A piecewise linear map between G-trees in the same deformation space is a quasi-isometry ([GL07, Remark 3.9]). Thus there is a constant C such that for all $x, y \in T$, $f(x) = f(y) \Rightarrow d(x, y) < C$, and the constant $\mathrm{Lip}(f)C$ has the property above.

Let $\phi \in \mathrm{Aut}(G)$. We say that a map $f : T \to T$ *represents* ϕ if it is ϕ-equivariant, i.e. for every $t \in T$ and $g \in G$ we have $f(g \cdot t) = \phi(g) \cdot T$.

We call $f : T \to T$ a *train track representative* of ϕ if it represents ϕ and is a train track map. A train track representative for $\Phi \in \mathrm{Out}(G)$ is a train track representative for some $\phi \in \Phi$.

Remark 2.1.6. For all $\Phi \in \mathrm{Out}(G)$, for all $T \in \mathcal{D}$, there exist a ϕ-equivariant map $f : T \to T$ with $\phi \in \Phi$ sending vertex to vertex and piecewise linear on edges. However the train track condition is restrictive.

A theorem by Bestvina and Handel in [BH92] states that every fully irreducible automorphism of F_N has a train track representative. It is not known whether an analogue is true for all GBS groups. Given an arbitrary representative for Φ one can apply a procedure to try to obtain a train track map but the termination has not been proved, unless $\mathcal{D}$ is finite dimensional ([Mei15]).

The *transition matrix* $A(f)$ associated to $f : T \to T$ is a matrix defined as follows. By assumption f sends vertex to vertex and edge to edge path. Suppose $E(T/G)$ contains n unoriented edges $e_1, \dots, e_n$. The matrix $A(f)$ is the square matrix of size n such that $A(f)_{i,j}$ is the number of edges in the orbit e_i which appear in $f(e_j)$, without taking orientation into account.

A subforest of a tree T is *proper* if it is not T. It is *essential* if it contains the axis of an element of G.

Lemma 2.1.7. *Let $\phi \in \mathrm{Aut}(G)$. The following assertions are equivalent:*

- *there exists a proper special factor $H < G$ such that the conjugacy class of H is invariant by ϕ*
- *there exists a tree T and a map $f : T \to T$ representing ϕ such that T contains a proper G-invariant essential subforest Y with $f(Y) \subset Y$.*

Proof. Suppose there exists a proper special factor $H < G$ such that H is invariant by ϕ. By definition of special factors there exists a collapse $\pi : T \to \bar{T}$ with $T \in \mathcal{D}$ such that H is the stabilizer of a vertex $v \in \bar{T}$. Then $Y := \pi^{-1}(G \cdot v)$ is a proper G-invariant subforest of T. Moreover it contains the axis of every loxodromic element of H so it is essential.

Now let us construct the map $f : T \to T$. The subgroup H is a GBS group and $\phi_{|H}$ is an automorphism of H. There exists a map $f_H : T_H \to T_H$ representing $\phi_{|H}$. Here is a construction of the map f_H. Let u be a vertex in T_H. In the collapse $T \to \bar{T}$, the vertex u is sent to v with $G_v = H$ so $G_u \subset H$. Since H is ϕ-invariant we also have $\phi(G_u) \subset H$. The point is that there exists a vertex $w \in T_H$ such that $\phi(G_u)$ fixes w: there exists $w \in T$ such that $\phi(G_u) \subset G_w$, and by replacing w by its projection on T_H we still have the inclusion.

Define $f_H(gu) = \phi(g)w$ for every $g \in G$. Since there are finitely many of G-orbits of vertices in T_H, one can repeat this procedure until f_H is fully defined on vertices of T_H and extend to edges by linearity.

Now we want to extend f_H to T. It suffices to define the image of one vertex in every orbit of vertices. Let v be a vertex in $T \setminus G \cdot H$. There exists a vertex $w \in T$ such that $\phi_{|H}(G_v) \subset G_w$. Define $f(v) = w$ and extend by equivariance on $G \cdot v$ by

$$f(gv) = \phi(g)f(v)$$

By repeating this for every orbit of vertices and extending linearly on edges, we define f such that $f_{|T_H} = f_H$. Therefore $f(Y) \subset Y$.

Conversely suppose there exists $T \in \mathcal{D}$ containing a proper G-invariant essential subforest Y. Let Y_0 be a connected component of Y containing the axis of a loxodromic element h. By collapsing the forest Y we obtain a tree $\bar{T}$ in which h fixes a point v. The stabilizer of v is a proper special factor H. There are finitely many orbits of vertices in $\bar{T}$ so there are finitely many conjugacy class of non-cyclic vertex stabilizers in $\bar{T}$. Since $f(Y) \subset Y$, f induces a map $\bar{f}$ on $\bar{T}$. The vertex v must be sent to a vertex with non-cyclic stabilizer. Thus there exists $n \geq 1, k \geq 1$ and $g \in G$ such that $f^{n+k}(v) = gf^k(v)$ so $\phi^{n+k}(H) \subset g\phi^k(H)g^{-1}$, i.e.

$$\phi^n(H) \subset g'Hg'^{-1}$$

with $g' = \phi^{-k}g$. By repeating this we obtain a sequence of decreasing special factors

$$H \supset g_1^{-1}\phi^n(H)g_1 \supset \cdots \supset g_N^{-1}\phi^{nN}(H)g_N \supset \ldots$$

By Corollary 1.4.10 of Chapter 1 the sequence must be stationary so H and $\phi^n(H)$ are actually conjugate. □

A square non-negative matrix A of size m is *irreducible* if for every $1 \leq i, j \leq m$ there exists $n \in \mathbb{N}$ such that $(A^n)_{i,j} > 0$. The matrix A is *primitive* if there exists $n > 0$ such that all coefficients of A^n are positive.

A well-known result about primitive and irreducible matrices is:

Theorem 2.1.8 (Perron-Frobenius)**.** *Let A be a non-negative primitive matrix with size $n \times n$.*

- *There exists a real eigenvalue $\lambda > 0$ (the* Perron-Frobenius eigenvalue*) such that for every other eigenvalue $\mu \neq \lambda$ we have $|\mu| < \lambda$.*
- *The eigenvectors for λ are unique up to scalar multiplication and there exists an eigenvector v for λ such that $v > 0$.*

If A is irreducible the same hold except that $|\mu| \leq \lambda$.

A proof of the theorem can be found in [Sen81, Theorem 1.1].

Lemma 2.1.9. *Suppose ϕ is a fully irreducible automorphism with a train track map $f : T \to T$. When ϕ is fully irreducible then there exists a collapse $\pi : T \to T'$ with $T' \in \mathcal{D}$ and a train track map $f' : T' \to T'$ with $\pi \circ f = f' \circ \pi$ such that $A(f')$ is an irreducible matrix.*

Moreover, if the Perron-Frobenius eigenvalue of $A(f')$ is greater than 1, then $A(f')$ is primitive, so there exists a power n such that $A(f')^n > 0$.

In that case f' is an *irreducible* train track map. If $A(f')$ is primitive then we call f *primitive*.

Proof. Suppose $A(f)$ is not irreducible. There exists a partition $I \cup J = \{1, \dots, n\}$ such that for every $i \in I$, $j \in J$, $A(f)_{i,j} = 0$. Then let Γ_J be the subgraph of $\Gamma = T/G$ spanned by edges e_j such that $j \in J$. Then the corresponding subforest $T_J \subset T$ is invariant by f. By Lemma 2.1.7, T_J is non essential since ϕ has no invariant conjugacy class of special factors.

Let $\pi : T \to \bar{T} = T/\sim_{T_J}$ be the collapse of the non essential invariant forest T_J. Since T_J is non essential $\bar{T} \in \mathcal{D}$. We have $f(T_J) \subset J$ so the map f induces a map $\bar{f}$ on $\bar{T}$ such that if π is the collapse map $T \to \bar{T}$ then $\pi \circ f = \bar{f} \circ \pi$.

If $A(\bar{f})$ is irreducible then we are done. Otherwise this process can be iterated by collapsing a non essential f-invariant forest of $\bar{T}$. At each step the number of orbits of edges in $\bar{T}$ decreases so this stops eventually and we obtain the map $f' : T' \to T'$ such that $A(f')$ is irreducible.

Let us check that the map $f' : T' \to T'$ is a train track map. Let π be the collapse map $T \to T'$. For every $k \in \mathbb{N}$ we have $\pi \circ f^k = f'^k \circ \pi$. If γ is a non-backtracking path in T then $\pi(\gamma)$ is also non-backtracking. For every edge $e' \in E(T')$, there is a unique edge $e \in \pi^{-1}(e)$ and $f'^k(e') = \pi \circ f^k(e)$. Since f is train track $f^k(e)$ is geodesic and so is $f'^k(e')$.

Let us prove that $A(f')$ is a primitive matrix. Since f' is a train track map, for every $n \in \mathbb{N}$, $A(f'^n) = A(f')^n$. As in [Sen81, Section 1.4] an irreducible matrix A admits a partition of indices such that A induces a permutation of the classes of the permutation. In terms of tree maps, it means that there exists $k \geq 1$ and a G-invariant partition $T_1 \sqcup \cdots \sqcup T_k$ of $E(T)$ such that f' induces a permutation of the subforests $T_i, i \in \{1, \dots, k\}$. Moreover $k = 1$ if and only if $A(f')$ is primitive.

By contradiction, suppose $k > 1$. There is a power f'^m such that for every $i \in \{1, \dots, k\}$ we have $f'^m(T_i) \subset T_i$. Lemma 2.1.7 implies that T_i must be a non-essential, thus collapsible, subforest. Its connected components must be uniformly bounded.

The Perron-Frobenius of $A(f')$ is greater than 1 so there is an edge e in T such that the diameter of $f^{mn}(e)$ is unbounded when n goes to infinity. There is a subforest T_i such that $e \in T_i$, and $f^{mn}(e)$ must also be contained in a connected component of T_i. This is a contradiction and it implies that $k = 1$ and $A(f')$ is primitive. □

Observe that the construction of $f' : T' \to T'$ starting from $f : T \to T$ is algorithmic.

Remark 2.1.10. Let f be a train track representative for ϕ (i.e. sending edges to legal paths but not necessarily at a uniform speed). Suppose $A(f)$ is irreducible. Let λ be the Perron-Frobenius eigenvalue for $A(f)$ and let v be the right Perron-Frobenius eigenvector normalized such that its coordinates add up to 1. The metric on T can be redefined such that $\text{len}(e_i)$ is the i-th coordinate of v. Then the Lipschitz constant of the new map $f : T \to T$ is uniform on edges and is equal to λ, so f is a metric train track map.

Because of Remark 2.1.10 all train track maps will be considered as metric train track maps. This point of view is not necessary but it gives some intuition on the behaviour of the iterate images of edges by f.

Remarks 2.1.11.

— If G is solvable then it has no special factor. Baumslag-Solitar groups have no fully irreducible automorphism: if $q = pn$ for some $n \in \mathbb{N}$, then $\text{BS}(p, q)$ has an $\text{Aut}(G)$-invariant conjugacy class of special factor ([Bou16]). For other Baumslag-Solitar groups or an amalgamated product $\mathbb{Z} *_{\mathbb{Z}} \mathbb{Z}$ all automorphisms have finite order in $\text{Out}(G)$ ([Lev07]). Thus we do not consider these groups when studying fully irreducible automorphisms.

— Suppose G is not solvable and is neither a Baumslag-Solitar group $\text{BS}(p, q)$ nor an amalgamated product $\mathbb{Z} *_{\mathbb{Z}} \mathbb{Z}$: then graphs of groups for G have more than one edge and have a proper subgraph representing a proper special factor. Let $f : T \to T$ be an irreducible train track map representing $\phi \in \text{Aut}(G)$. If the Perron-Frobenius eigenvalue of $A(f)$ is 1 then $T = T \cdot \phi$ and f is an isometry. This implies that f preserves a subgraph of T/G so ϕ is not fully irreducible.

Corollary 2.1.12. *Suppose G is a non-solvable Baumslag-Solitar group which is neither a Baumslag-Solitar group nor an amalgamated product $\mathbb{Z} *_{\mathbb{Z}} \mathbb{Z}$. Let $\Phi \in \text{Out}(G)$ with a train track representative.*

There is an algorithm which takes a train track representative $f : T \to T$ for Φ and either computes a primitive train track representative for Φ or gives a proof that Φ is reducible.

Proof. The construction of an irreducible representative for Φ follows Lemma 2.1.9. Lemma 2.1.9 also states that if Φ is fully irreducible and the Perron-Frobenius eigenvalue is strictly greater than 1 then the transition matrix is also primitive. By Remark 2.1.11 2. if the Perron-Frobenius eigenvalue is 1 then Φ is reducible. In that case the matrix is not primitive. Since there exists a power n depending on the size of the matrix A such that if A is primitive then $A^n > 0$, one can algorithmically test the primitivity of the transition matrix. □

2.1.3 The stable lamination associated to a train track map

Let T be a G-tree in $\mathcal{D}$. Let ∂T be the space of ends of T, i.e. the set of equivalence classes of infinite geodesic rays of T, where rays ρ, ρ' are equivalent if the Hausdorff distance $d(\rho, \rho')$ is finite. We endow it with the standard topology: a basis of neighbourhoods for $\xi \in \partial T$ is $\{\mathcal{V}_x, x \in V(T)\}$ where $\mathcal{V}_x$ is the connected component of $T \setminus \{x\}$ containing a ray for ξ.

For any other tree $T' \in \mathcal{D}$ there exists a G-equivariant quasi-isometry $f : T \to T'$ which induces a G-equivariant homeomorphism $\partial f : \partial T \to \partial T'$. The homeomorphism ∂f does not depend on the choice of f so there is a canonical identification of $\partial T'$ with ∂T for any $T' \in \mathcal{D}$.

A *lamination* of G is a closed, symmetric, G-invariant subset of $\partial T \times \partial T \setminus \Delta$, for any tree $T \in \mathcal{D}$, where Δ is the diagonal. As discussed above, its definition does not depend on the tree T.

There is a canonical action of $\mathrm{Aut}(G)$ on ∂T and thus on the set of laminations: let $T \in \mathcal{D}$ and $\phi \in \mathrm{Aut}(G)$ and let $f : T \to T$ be a representative for ϕ. It induces a homeomorphism of ∂T which does not depend on the choice of f.

Since laminations are G-invariant, the action of $\mathrm{Aut}(G)$ on ∂T yields an action of $\mathrm{Out}(G)$ on the set of laminations.

Let Λ be a lamination. For $T \in \mathcal{D}$, the *realization* Λ_T of Λ in T is the G-invariant set of unoriented bi-infinite geodesics whose endpoints belong to $\Lambda \subset \partial T \times \partial T \setminus \Delta$. The geodesics of the realization are called *leaves* of Λ. If $T' \in \mathcal{D}$ and $f : T \to T'$ is a quasi-isometry, then for any leaf $\lambda \in \Lambda_T$, the geodesic obtained by tightening $f(\lambda)$ is a leaf of

$\Lambda_{T'}$. Conversely all leaves of $\Lambda_{T'}$ are obtained that way.

A *leaf segment* of Λ_T is a segment of T contained in a leaf in T.

We now introduce the *stable lamination* of an automorphism with a irreducible train track.

Definition 2.1.13. Let $\Phi \in \mathrm{Out}(G)$ and $f : T \to T$ be an irreducible train track representative for Φ. The *stable lamination* Λ^+ is defined by its realization in T. A bi-infinite geodesic λ lies in Λ_T^+ if and only if for every leaf segment $\sigma \subset \lambda$, there exists an edge $e \in E(T)$ and $n \in \mathbb{N}$ such that $\sigma \subset f^n(e)$.

Remark 2.1.14. Since all edges are legal, all leaf segments are also legal, so there is no cancellation in leaves in T when applying f. Thus leaves of Λ_T^+ are legal. Viewing f as a metric train track map, they are also uniformly stretched by the factor $\mathrm{Lip}(f)$. The set Λ_T^+ is stable by f. It implies that Λ^+ is stable by Φ.

Remark 2.1.15. If $\Phi \in \mathrm{Out}(G)$ admits an irreducible train track representative $f : T \to T$ and thus a stable lamination Λ^+ can be defined, then Φ^n admits $f^n : T \to T$ as a train track representative. Then the stable lamination associated to f^m is equal to Λ^+. Indeed, a subsegment of $f^k(e)$ is also a subsegment of $f^{nm}(e)$ for some $m \in \mathbb{N}$: since for N big enough, $f^N(e)$ contains a translate of e, then for m big enough $f^{nm-k}(e)$ contains a translate of e and $f^{nm}(e)$ contains a translate of $f^k(e)$.

Definition 2.1.16. A lamination Λ is *minimal* with respect to a tree T if it satisfies the following condition: $\forall \lambda, \lambda' \in \Lambda_T$, $\forall I$ leaf segment in λ, $\exists g \in G$ such that $gI \subset \lambda'$. In other words, all leaves of a minimal lamination have the same leaf segments up to the action of G.

Remark 2.1.17. Let $T, S \in \mathcal{D}$. A lamination Λ is minimal with respect to T if and only if it is minimal with respect to S. Indeed let $f : T \to S$ be a G-equivariant quasi-isometry. There exists $C > 0$ depending on f such that the image of any geodesic path in T is in the C-neighbourhood of a geodesic path in S.

Let λ_S, λ'_S be leaves of Λ_S. Let λ_T, λ'_T be the corresponding leaves in Λ_T. Let $I_S \subset \lambda_S$ be a leaf segment. There exists $I_T \subset \lambda_T$ such that $f(I_T)$ contains a C-neighbourhood of I_S in λ_S. There also exists $g \in G$ such that $gI_T \subset \lambda'_T$. The segment $f(gI_T) = gf(I_T)$ is in a C-neighbourhood of λ'_S so its central part containing gI_S is contained in λ'_S.

Lemma 2.1.18. *The stable lamination associated to an automorphism with an irreducible train track representative is minimal.*

A proof can be found in [BFH97, Lemma 1.2] for free groups, and it can be adapted to the case of GBS groups.

Definition 2.1.19. A leaf ℓ of the realization of a lamination in a tree is *quasi-periodic* if for every $l > 0$ there exists $C_l > 0$ such that if σ is a leaf segment of ℓ of length at most l then any leaf segment of ℓ of length at least C_l contains a translate of σ.

Lemma 2.1.20. *Let $\Phi \in \mathrm{Out}(G)$ with a train track $f : T \to T$. Let ℓ be a leaf of Λ^+_T. The leaf ℓ is quasi-periodic.*

Proof. Let $l > 0$. Let $e \in E(T)$. Since the stable lamination is minimal, there exists $k \in \mathbb{N}$ depending on l such that $f^k(e)$ contains an edge path in every orbit of edge paths with length at most l crossed by leaves of the lamination.

There exists $n \in \mathbb{N}$ such that for every $e' \in E(T)$ and $m \geq n$, the path $f^m(e')$ contains an edge in the orbit of e.

Let $C_l := \max_{e \in E(T)} \mathrm{len}(f^{n+k}(e'))$. Let $\sigma \subset \ell$ be a leaf segment with length at least C_l. There exists $m \in \mathbb{N}$ and $e' \in E(T)$ such that $\sigma \subset f^m(e')$. Moreover $m \geq n + k$. The path $f^{m-k}(e')$ contains an edge in the orbit of e so $f^m(e')$ contains a subpath in the orbit of $f^k(e)$, which itself contains a path in every orbit of edge paths with length at most l. □

Let us fix an automorphism $\Phi \in \mathrm{Out}(G)$, and suppose it admits a train track representative and thus a stable lamination Λ^+. Let $S \in \mathcal{D}$, not necessarily the train track representative. Let v be a vertex in S. Let λ be a leaf in the stable lamination Λ^+_S. The *Whitehead graph* of the leaf λ at the vertex v is the graph $\mathrm{Wh}_S(\lambda, v)$ such that

— vertices are edges of S with origin v

— there is an edge $e - e'$ if there exists $g \in G$ such that $\{e, e'\}$ is a turn in $g\lambda$

The vertex stabilizer G_v acts naturally on $\mathrm{Wh}_S(\lambda, v)$. The stable lamination is minimal so all leaves of the stable lamination have the same subsegments up to translation. In particular they have the same turns up to translation so for any leaf λ' in Λ^+ we have $\mathrm{Wh}_S(\lambda', v) = \mathrm{Wh}_S(\lambda, v)$. We may as well define the Whitehead graph of the stable lamination at vertex v by $\mathrm{Wh}_S(\Lambda^+, v) = \mathrm{Wh}_S(\lambda, v)$ for any leaf λ.

Lemma 2.1.21. *Let $\Phi \in \mathrm{Out}(G)$ be an automorphism with an irreducible train track representative $f : T \to T$ and associated stable lamination Λ^+. The Whitehead graphs $\mathrm{Wh}_T(\Lambda^+, v)$ can be computed algorithmically.*

Proof. In order to compute Whitehead graphs we need to find all orbits of turns taken by the lamination.

Let $e \in E(T)$. Since f is irreducible, every leaf segment of Λ_T^+ is a subsegment of a translate of $f^n(e)$ for some $n \in \mathbb{N}$. Therefore all turns taken by Λ_T^+ appear in $f^n(e)$ for some $n \in \mathbb{N}$.

There are two ways a turn can arise in some $f^n(e)$. If $\{e_1, e_2\}$ is a turn in $f^n(e)$ then either it is a turn in $f(e')$ for some $e' \in E(T)$ (first type), or it is the turn between $f^k(e_1')$ and $f^k(e_2')$ where $\{e_1', e_2'\}$ is a turn of the first type (second type).

There exists a power f^n depending only on the size of $A(f)$ such that for every edge e of T, $f^n(e)$ crosses every orbit of edge in T. Therefore all turns of the first type appear in $f^{n+1}(e)$. Let K be the number of orbits of turns in T, which is finite. Then all turns of the second type appear in $f^{n+1+k}(e)$ for some $k \leq K$. This gives a bound on the number of iterations needed to find all turns. □

The following results may apply to more general laminations. Let Λ be a lamination. We say that a finitely generated subgroup $H \subset G$ containing a hyperbolic element *carries* Λ if there exists $S \in \mathcal{D}$ such that any leaf in Λ_S is contained in a translate of the minimal subtree for H in S. The condition only depends on the conjugacy class of H. Moreover it does not depend on a choice of S:

Lemma 2.1.22. *If a finitely generated subgroup H carries Λ then in any $S \in \mathcal{D}$, any leaf of Λ_S is contained in a translate of the minimal subtree S_H.*

Proof. Let $S \in \mathcal{D}$. A bi-infinite geodesic λ is contained in S_H if and only if it lies in a bounded neighbourhood of S_H.

Since H carries Λ, there exists a tree $S' \in \mathcal{D}$ such that all leaves of $\Lambda_{S'}$ are contained in S'_H.

Let $f : S' \to S$ be a G-equivariant map. It is a quasi-isometry. Let λ be a leaf of $\Lambda_{S'}$. Up to translating it we may suppose that it is contained in S'_H. Then $f(\lambda)$ is contained in $f(S'_H)$. By minimality the tree S_H is contained in $f(S'_H)$. The diameter of S'_H/H is bounded, also by minimality, so the diameter of $f(S'_H)/H$ is also bounded. Therefore $f(S'_H)$ must be contained in a c-neighbourhood of S_H for some $c > 0$.

The image $f(\lambda)$ must then be contained in a c-neighbourhood of S_H. The tightened leaf $[f(\lambda)]$ is a geodesic contained in $f(\lambda)$ so it is contained in S_H. □

For minimal laminations such as the stable lamination of an automorphism, we have the following:

Proposition 2.1.23. *Let Λ be a non-empty minimal lamination. Let H be a special factor of G. Then Λ is carried by H if and only if there exists $T \in \mathcal{D}$ such that there exists a leaf of Λ_T contained in the minimal subtree for H in T.*

Proof. The direct implication is immediate. Conversely, by Lemma 2.1.22, if there exists $T \in \mathcal{D}$ and a leaf $\lambda \in \Lambda_T$ contained in the minimal subtree T_H, then for all $T' \in \mathcal{D}$, the realization of λ in $\Lambda_{T'}$ is also contained in T'_H.

We choose $T' \in \mathcal{D}$ such that T'_H and its translates are disjoint: for all $g \in G$, $gT'_H \cap T_H \neq \varnothing \Rightarrow g \in H$. The leaf λ is contained in T'_H. If λ' is another leaf, its segments are translates of segments of λ. This implies that λ' all edges crossed by λ' belong to $G \cdot T'_H$. As translates of T'_H are disjoint λ' must be contained in a single translate of T'_H. □

Lemma 2.1.24. *Let A be a special factor for G. Let $T \in \mathcal{D}$. Let T_A be the minimal subtree for A. There exists $C > 0$ such that for all $g \in G \setminus A$ we have* $\operatorname{diam}(T_A \cap g \cdot T_A) \leq C$.

Proof. By definition of a special factor, there exists a tree $S \in \mathcal{D}^{\mathcal{A}}$ such that the translates of the minimal subtree S_A are disjoint.

There exists a G-equivariant application $f : S \to T$ which is a quasi-isometry and such that $T_A \subset f(S_A)$. There exists a constant C depending on f such that if two subtrees are disjoint, the diametre of their intersection is bounded by C. □

We deduce the following, which can be applied to the stable lamination of an automorphism in our context:

Proposition 2.1.25. *Let Λ be a minimal lamination. There exists a unique special factor relative to $\mathcal{D}$ and minimal for inclusion which carries Λ.*

Proof. The existence is a consequence of the descending chain condition given in Corollary 1.4.10 of Chapter 1. For uniqueness we need to show that if a leaf of Λ is contained in the minimal subtree for two special factors A and A' then it is also contained in the minimal subtree for $A \cap A'$.

It suffices to prove that the intersection $T_A \cap T_{A'}$ is contained in a bounded neighbourhood of $T_{A \cap A'}$. By Lemma 2.1.24 there exists $C > 0$ such that for every $g \in G$, $\operatorname{diam}(T_A \cap gT_A) \geq C \Rightarrow g \in A$ and $\operatorname{diam}(T_{A'} \cap gT_{A'}) \geq C \Rightarrow g \in A'$.

There exists $J > 0$ depending on C and T such that every segment σ of T with length at least J contains at least two segments of length C in the same G-orbit and with same orientation in σ.

Let $\sigma \subset T_A \cap T_{A'}$ be a segment with length at least J. There exists a segment $I \subset \sigma$ with length C and $g \in G$ such that $gI \subset \sigma$. Then the diameter of $T_A \cap gT_A$ is greater than C so $g \in A$. Similarly $g \in A'$. The axis of g is contained in $T_{A \cap A'}$. Thus $T_A \cap T_{A'}$ is contained in a J-neighbourhood of $T_{A \cap A'}$. □

2.2 Whitehead graphs of the lamination and reducibility

In order to decide whether an automorphism with a train track representative is fully irreducible, the Whitehead graphs of the stable lamination give important information. In order to state the main theorem we need to introduce first pseudo-periodic conjugacy classes, which will be developed further in Section 2.3.

Definition 2.2.1. Let $\Phi \in \mathrm{Out}(G)$ and let $\phi \in \Phi$.

The conjugacy class of an element $g \in G$ is *pseudo-periodic* for Φ if $\|\phi^n(g)\|_T$ is bounded.

A fact worth mentioning is that the minimal special factor containing a pseudo-periodic element is periodic. We prove this in Section 2.3.

The aim of this section is to prove:

Theorem 2.2.2. *Let Φ be an automorphism of G with a train track $f : T \to T$ with no simple pseudo-periodic conjugacy class. Then Φ is fully irreducible if and only if for every $v \in T$, the Whitehead graph $\mathrm{Wh}_T(\Lambda_f^+, v)$ is connected.*

Propositions 2.2.3 and 2.2.6 are GBS equivalents for analogue results by Kapovich ([Kap14, Proposition 4.1, Proposition 4.2]). Along with Corollary 2.2.7 they prove Theorem 2.2.2.

Proposition 2.2.3. *Let $\Phi \in \mathrm{Aut}(G)$ be an automorphism with an irreducible train track representative $f : T \to T$ and an associated stable lamination Λ^+.*

Let $S \in \mathcal{D}$ be any tree. Let $f' : S \to S$ be a representative for Φ. If there exists $v \in S$ such that the Whitehead graph $\mathrm{Wh}_S(v, \Lambda^+)$ is not connected and such that the stabilizer of some of its connected components is in $\mathcal{A}$, then Φ is reducible in $\mathcal{D}^{\mathcal{A}}$.

Remark 2.2.4. If $\mathcal{D}^{\mathcal{A}} = \mathcal{D}$ the condition on stabilizers of connected components is always true.

Lemma 2.2.5. *Let Φ be an outer automorphism having an irreducible train track representative. Let S be a G-tree. If there exists a special factor whose minimal subtree in S contains a leaf of the stable lamination, then Φ is reducible.*

Proof. Let λ be a leaf of the stable lamination. There exists a unique minimal special factor H carrying the leaf λ (Proposition 2.1.25). By Proposition 2.1.23, every leaf of the lamination is contained by a translate of H. Let $\phi \in \Phi$. Since the stable lamination is ϕ-invariant, all leaves are also carried by translates of $\phi(H)$. By minimality there is an element h such that we have $H \subset h\phi(H)h^{-1}$. It follows that $H = h\phi(H)h^{-1}$. Otherwise, construct a decreasing sequence

$$\cdots \subset h_n^{-1}\phi^{-n}(H)h^n \subset \cdots \subset H$$

and by Chapter 1, Corollary 1.4.10 this sequence is stationary. This implies that H is conjugate to $\phi(H)$. □

Proof of Proposition 2.2.3.. Let $\phi \in \mathrm{Aut}(G)$ with a train track representative $f : T \to T$. Let $f' : S \to S$ be a representative for ϕ. Suppose there exists a vertex v such that the Whitehead graph $\mathrm{Wh}_S(v, \Lambda^+)$ is non-connected with at least one connected component with stabilizer in $\mathcal{A}$. We will construct a tree S' such that there exists a collapse $S' \to S$, and such that the lifts of the leaves of Λ^+ in S' avoid an orbit of edges.

Let $C_1, \ldots, C_k$ be the connected components of $\mathrm{Wh}_S(v, \Lambda^+)$.

Let E be the star with k edges $e_1, \ldots, e_k$. Replace v by E by attaching the end of e_i to edges in C_i. By extending this construction by G-equivariance, we define an expansion at vertex v.

The edge e_i has the same stabilizer as the corresponding connected component C_i. It is possible that the construction created an edge with stabilizer not in $\mathcal{A}$. Collapse these edges. By assumption, at least one edge in E has stabilizer in $\mathcal{A}$, so not all edges are collapsed.

We obtain a tree S' with a collapse map $\pi : S' \to S$. The leaves of the stable lamination can be lifted in S'.

By construction leaves in S' do not cross the edges which are collapsed by π. Thus there is a special factor H which carries a leaf of the lamination. By Lemma 2.2.5, ϕ is reducible. □

Proposition 2.2.6. *Let $\Phi \in \mathrm{Out}(G)$. Let $f : T \to T$ be a train track representative for Φ. Suppose that the incidence matrix $A(f)$ is primitive. Let Λ^+ be the stable lamination.*

Then if for every $v \in T$, the Whitehead graph $\mathrm{Wh}_T(\Lambda^+, v)$ is either connected, or disconnected such that no connected component has a stabilizer in $\mathcal{A}$, then no leaf of the lamination is carried by a proper special factor.

Proof of Proposition 2.2.6.. First let us prove the proposition under simplified assumptions: we suppose that all Whitehead graphs are connected.

By contradiction, let Φ, f, T be as in the proposition and such that all Whitehead graphs of the stable lamination are connected. Let $\phi \in \Phi$ be such that f is ϕ-equivariant. Let A be a special factor such that the leaf λ is contained in the minimal subtree T_A of A in T. By Proposition 2.1.23 this implies that A carries the stable lamination and that every leaf of it is contained in a translate of T_A.

Let e be an edge in T_A. Let e' be any edge of T. Up to reversing the orientation of these edges, there exists a geodesic path $e = e_0 \dots e_n = e'$ (see Figure 2.2 for what follows). Turns in this path are not necessarily crossed by any leaf of Λ^+. However, Whitehead graphs are connected. At every turn $\bar{e}_i, e_{i+1}$ at vertex v_i, there exist edges $\bar{e}_i = \epsilon_0, \dots, \epsilon_k = e_{i+1}$ with origin v_i such that $\epsilon_j, \epsilon_{j+1}$ is a turn crossed by a leaf. Thus $\bar{\epsilon}_j \epsilon_{j+1}$ is a leaf segment.

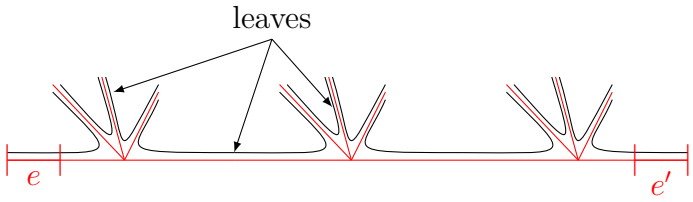

Figure 2.2 – Cover the path between e and e' with overlapping leaf segments.

This yields a sequence of leaf segments which cover the path between e and e' and such that two consecutive segments overlap over at least the length of an edge.

Since leaf segments are legal, they are all stretched uniformly by f. There is a power m of f such that the image of any edge by f^m is strictly longer than constant C given by Lemma 2.1.24.

If we apply f^m to the leaf segments above, we obtain longer leaf segments such that two consecutive segments overlap by a length greater than C. Since all leaf segments are contained in translates of T_A, they must be contained in the same translate $h_{e'}T_A$. In particular $f^m(e)$ and $f^m(e')$ are both contained in $h_{e'}T_A$, by Lemma 2.1.24.

Since $f^m(e)$ is longer than C and $f^m(e) \subset h_{e'}T_A$ for any $e' \in E(T)$, the element $h_{e'} \in G$ actually does not depend on e' and we denote it by h.

Therefore there exists $m \in \mathbb{N}$ and $h \in G$ such that $f^m(e') \subset hT_A$ for all $e' \in E(T)$. This implies that $f^m(T) \subset T_A$, which gives the desired contradiction since f^m is surjective by minimality of the action.

Now suppose that Whitehead graphs may be disconnected, though stabilizers of connected components do not belong to $\mathcal{A}$. Let A be a special factor such that T_A contains a leaf λ of the stable lamination.

Let $\mathcal{R}$ be the smallest G-equivariant equivalence relation on $E(T)$ such that two edges in a turn crossed by a leaf of the lamination belong to the same equivalence class. It defines a partition of $E(T)$. If $e\mathcal{R}e'$ then all edges on the geodesic path between e and e' also belong to the same equivalence class. Each equivalence class spans a connected subtree of T. When all Whitehead graphs are connected, there is a single equivalence class.

Now we prove that there exists $m \in \mathbb{N}$ such that $f^m(S_{i_0})$ is contained in a translate of T_A. The argument is the same as in the simple case of the proof: take $e \in T_A \cap S$ on the leaf λ. There is $m \in \mathbb{N}$ depending uniquely on T, A, f such that for any $e' \in S$ the images $f^m(e)$ and $f^m(e')$ are in the same translate of T_A. This works well since the restriction of any Whitehead graph to S_{i_0} is connected. Then we have $f^m(S_{i_0}) \subset gT_A$ for some $g \in G$.

As $f^m(S) \subset gT_A$ and $f^m(S_{i_0})$ is unbounded, Lemma 2.1.24 implies $\phi^m(\operatorname{Stab} S) \subset \operatorname{Stab} f^m(S_{i_0}) \subset gAg^{-1}$, so $\operatorname{Stab}(S) \subset \phi^{-m}(gAg^{-1})$.

Let $\{S_i\}_{i\in I}$ be the set of subtrees of T corresponding to equivalence classes of $\mathcal{R}$. The intersection $S_i \cap S_j$ for $i \neq j$ is either empty or a single vertex. This means that $\{S_i\}_{i\in I}$ is a *transverse covering* of T (see [Gui04, Def. 4.6]).

There exists $i_0 \in I$ such that λ is contained in S_{i_0}. Since the transition matrix of the train track is primitive, λ crosses all orbits of edges in T so all subtrees S_i are translates of S_{i_0}. Construct $\hat{T}$ (the *skeleton of the transverse covering*, [Gui04, Def. 4.8]) as follows. It is a bipartite graph. Vertices of $\hat{T}$ are s_h with $h \in G/\operatorname{Stab} S_{i_0}$ which correspond to translates of S_{i_0} and x_v for $v \in T$ such that v is the intersection of two distinct translates of S_{i_0}. There is an edge $s_h - x_v$ if $v \in hS$. This yields a G-tree with vertex stabilizers of x_v elliptic in $\mathcal{D}$ and vertex stabilizers of s_h conjugate to $\operatorname{Stab} S$ [1]. The edge group of $s_h - x_v$ is the stabilizer of the component of $\operatorname{Wh}_T(\Lambda^+, v)$ which intersects hS_{i_0}, so it is not in $\mathcal{A}$.

1. This tree may also be obtained by replacing vertices of T by stars on the connected components of their Whitehead graph, as in the proof of Proposition 2.2.3, then collapsing copies of S_{i_0}.

Construct a map $\hat{T} \to R$ where R is a collapse of an element of $\mathcal{D}^{\mathcal{A}}$ and has a vertex stabilizer $\phi^{-m}(A)$. Map the vertex s_h of $\hat{T}$ for $h \in G$ to the vertex of R with stabilizer $h\phi^{-m}(A)h^{-1}$ and map every vertex x_v to a vertex of R whose stabilizer contains G_v. All edge groups in R are in $\mathcal{A}$ while no edge group in $\hat{T}$ is in $\mathcal{A}$. Thus every edge of $\hat{T}$ is sent to a point in R. By continuity the image of $\hat{T}$ is a single point in R and is G-invariant. By minimality of the action on R, the tree R must be a single point and $\phi^{-m}A = G$ so $A = G$, which is a contradiction. □

Corollary 2.2.7. *If $\Phi \in \mathrm{Out}(G)$ satisfies the hypotheses of Proposition 2.2.6, and no simple element in G is pseudo-periodic for Φ, then Φ is fully irreducible.*

Proof. By contradiction, let Φ be an automorphism satisfying the hypotheses of the corollary. Let $\phi \in \Phi$.

Let H be a special factor of G. Up to replacing Φ by Φ^k and to choosing $\phi \in \Phi^k$, we may assume that $\phi(H) = H$.

Let T_H be the minimal subtree of H. Let $h \in H$ be a loxodromic element whose axis is contained in T_H, so h is simple. By assumption $\|\phi^n(h)\|$ goes to infinity.

We will find a contradiction to Proposition 2.2.6 by finding a leaf of Λ^+ contained in T_H.

Let $e_1, \dots, e_m$ be the edges of a fundamental domain for h. Some fundamental domain of the axis of $\phi^n(h)$ can be written as a concatenation of (maybe empty) sets J_i^n where $J_i^n \subset f^n(e_i)$ for $1 \leq i \leq m$.

The non-degenerate J_i^n are leaf segments of Λ^+, some of them may be empty. Since $\|\phi^n(h)\|_T$ goes to infinity, at least one of the J_i^n must also go to infinity when n goes to infinity. Besides we have $J_i^n \subset T_{\phi(H)} = T_H$ for all $n \in \mathbb{N}$.

Suppose that $(J_{i_0}^n)_{n\in\mathbb{N}}$ tends to infinity. The segments $J_{i_0}^n$ are arbitrarily long leaf segments, all contained in T_H. Let λ be a leaf of Λ^+. Then every leaf segment of λ is contained in a translate of $J_{i_0}^n$ for some $n \in \mathbb{N}$.

By Lemma 2.1.24 there exists C such that if the intersection of two translates of T_H has diameter greater than C, then the translates are equal.

Since λ is quasi-periodic there exists $L > 0$ such that for every leaf segment σ of length $3C$, every leaf segment $\gamma \subset \lambda$ with $\mathrm{len}(\gamma) \geq L$ contains a translate of σ.

There exists i_0 and $n \in \mathbb{N}$ such that $\mathrm{len}(J_{i_0}^n) \geq L$. By minimality of Λ^+, the leaf λ contains a translate of $J_{i_0}^n$ so every leaf segment $\sigma \subset \lambda$ with $\mathrm{len}(\sigma) = 3C$ is a translate of a subsegment of $J_{i_0}^n \subset T_H$. Tile λ with such leaf segments of length C such that two

consecutive tiles overlap on more than C. Each segment of the tiling is contained in a translate of T_H and by Lemma 2.1.24 all these translates must be equal. Thus $\lambda \subset gT_H$ for some $g \in G$ and Proposition 2.1.23 implies that H carries Λ. □

2.3 Nielsen paths and pseudo-periodic elements

The aim of the present section is to define and describe *pseudo-periodic* conjugacy classes of an outer automorphism $\Phi \in \mathrm{Out}(G)$. They are analogues of periodic conjugacy classes for an automorphism of a free group.

Pseudo-periodicity is a weaker notion than periodicity: the conjugacy class of $g \in G$ is pseudo-periodic if the translation length $\|\phi^n(g)\|_T$ is bounded for some, equivalently any, tree $T \in \mathcal{D}$, and some $\phi \in \Phi$. While it automatically implies that the conjugacy class of g is periodic in the free group case, it is not always true for GBS groups.

An outer automorphism $\Phi \in \mathrm{Out}(G)$ is *pseudo-atoroidal* if it has no pseudo-periodic conjugacy classes.

The point is that when there exists a *simple* element $g \in G$ whose conjugacy class is pseudo-periodic, then the conjugacy class of the unique proper minimal special factor containing g is Φ-periodic. In Section 2.2 we gave a criterion for reducibility, such that automorphisms satisfying the criterion are irreducible if and only if they do not have any pseudo-periodic conjugacy classes.

For algorithmic purposes we need to find the pseudo-periodic conjugacy classes for an automorphism $\Phi \in \mathrm{Out}(G)$. We will see that like for free groups, there is a strong link between pseudo-periodic conjugacy classes and *Nielsen paths*, which are some periodic paths in a train track representative for Φ. Finding Nielsen paths in a train track representative will be our first goal. Then we establish the link with pseudo-periodic conjugacy classes.

2.3.1 Computation of periodic indivisible Nielsen paths

Like in the free group case, we can define Nielsen paths associated to a train track map. They are a useful tool to understand periodic elements of an automorphism. In the free group Nielsen paths are usually defined in the quotient graph but here we will define them in the tree instead, since in GBS graphs some non-degenerate turns cannot be seen in the quotient graph. Some example of references include [BH92], [BFH97].

For any path α in a tree T we define $[\alpha]$ as the path obtained by tightening α while

keeping the same endpoints. Equivalently, it is the geodesic between the endpoints of α. If the endpoint of α is the initial point of β then $\alpha \cdot \beta$ is the concatenation of α and β. It is a *tight* concatenation if α and β are geodesics and $\alpha \cdot \beta$ is a geodesic. A tight concatenation is a *legal* concatenation if the turn at the concatenation point is legal, otherwise it is *illegal*.

Let $f : T \to T$ be an irreducible train track representative for $\Phi \in \mathrm{Out}(G)$. A *Nielsen path* is a tight path $\gamma \subset T$ such that there exists $g \in G$ such that $f(\gamma)$ and $g\gamma$ are homotopic relative endpoints, i.e. $g^{-1} \cdot f$ fixes both endpoints of γ.

A *periodic Nielsen path* is a tight path $\gamma \subset T$ such that there exists $n \geq 1$ and $g \in G$ such that $f^n(\gamma)$ and $g\gamma$ have the same endpoints. The minimal nonzero integer n is the *period* of γ.

Remark 2.3.1. A tight concatenation of periodic Nielsen paths $\gamma_1 \cdot \gamma_2$ might not be a periodic Nielsen path but it is pre-periodic: there exists $k \in \mathbb{N}$ such that $[f^k(\gamma_1 \cdot \gamma_2)]$ is a periodic Nielsen path. This is a consequence of the local finiteness of trees: $[f^k(\gamma_1 \cdot \gamma_2)]$ has bounded length. Since the tree is locally finite and the orbits of the endpoints are periodic, there are finitely many possibilities for the orbit of $[f^k(\gamma_1 \cdot \gamma_2)]$, hence its periodicity.

In general the endpoints of a Nielsen path are not vertices of the tree but rather lie in the interior of some edges.

An *indivisible Nielsen path* (INP) is a Nielsen path which cannot be written as the tight concatenation of two shorter Nielsen paths. A *periodic indivisible Nielsen path* (pINP) is a periodic Nielsen path which cannot be written as the concatenation of two shorter periodic Nielsen paths. Periodic INPs are defined as INPs for a certain power of f. Just as for periodic Nielsen path one can define the *period* of a periodic INP.

Lemma 2.3.2 is a GBS version of well-known facts for pINPs in the free group case.

Lemma 2.3.2. *Let $f : T \to T$ be a train track map.*

(i) Let γ be a periodic Nielsen path of period n. There exists a unique decomposition of γ as a tight concatenation of periodic indivisible Nielsen paths whose period divides n.

(ii) If γ is a tight concatenation of periodic indivisible Nielsen paths $\eta_1 \cdot \cdots \cdot \eta_k$ and η is a pINP such that $\eta \subset \gamma$ then there exists $i \in \{1, \ldots, k\}$ such that $\eta = \eta_i$. In particular the decomposition of a periodic Nielsen path into pINPs is unique.

(iii) A periodic indivisible Nielsen path contains a unique illegal turn. It is a tight concatenation $\alpha \cdot \beta$ where α, β are legal subpaths.

Proof. Let γ be a periodic Nielsen path. Let $g \in G$ and $n \in \mathbb{N}$ such that $g \cdot f^n$ fixes the endpoints of γ.

Let $\gamma_1, \dots, \gamma_k$ be the maximal legal subsegments such that $\gamma = \gamma_1 \cdot \dots \cdot \gamma_k$. The path γ cannot be legal since the lengths of its images by f^n is bounded so $k \geq 2$.

We have $\gamma_i \subset g \cdot f^n(\gamma_i)$ for all $1 \leq i \leq k$. Indeed $f^n(\gamma_i)$ is a legal segment, so it can only intersect a unique $g^{-1}\gamma_j$, otherwise it would cross an illegal turn. Since $f^n(\gamma)$ covers $g^{-1}\gamma$ and $[f^n(\gamma)]$ has exactly the same number of illegal turns as γ, f^n induces a permutation of the γ_i for $1 \leq i \leq k$. By continuity of f^n the permutation must be trivial.

Up to taking a multiple of n we may suppose that there is cancellation at every illegal turn when applying f^n. For $2 \leq i \leq k-1$, γ_i lies in the interior of $g \cdot f^n(\gamma_i)$. Since f^n stretches legal segments uniformly this implies that there is a unique fixed point v_i for $g \cdot f^n$ in the interior of γ_i. The vertices v_i cut γ into $k-1$ shorter periodic Nielsen paths, each of which has a unique illegal turn.

Since periodic indivisible Nielsen paths cannot be cut into smaller paths, they have at most one illegal turn, and at least one since their length does not grow exponentially with f. This proves (iii). Conversely, if a periodic Nielsen path has a unique illegal turn, it can be written as a concatenation of legal paths $\alpha \cdot \beta$. Up to replacing f by gf^n, the first point of α and the last point of β are fixed by f and all other points in $\alpha \cdot \beta$ escape exponentially when iterating f. Therefore pINPs are exactly periodic Nielsen paths with one illegal turn. For a general periodic Nielsen path γ, we found a decomposition into periodic Nielsen paths with one illegal turn each, also pINPs so (i) is proved.

Let us prove (ii). Now we do not assume that γ is a periodic Nielsen path, only that it is a concatenation of pINPs. We just proved that the former implies the latter. Observe that if α, β are legal paths and if $\eta = [x, y]$ is a pINP with $x \in \alpha$, $y \in \beta$ then for any $a, b \in \alpha \cdot \beta$ such that $\{a, b\} \neq \{x, y\}$, either $d(f^n(a), f^n(b)) \to \infty$ or there exists $n \in \mathbb{N}$ such that $f^n(a) = f^n(b)$.

Let $\eta = [x, y]$ be a pINP in γ. There exists a unique illegal turn in η and it is also the illegal turn in a pINP $\gamma' = [x', y']$ which appears in a decomposition of γ.

Let α, β be the maximal legal subsegments in which x, y lie. Since γ' has a unique illegal turn, x', y' lie in $\alpha \cdot \beta$. Since γ' is a pINP the sequence $d(f^n(x'), f^n(y'))$ is bounded and positive. The claim implies that $\{x', y'\} = \{x, y\}$.

This proves the uniqueness of the decomposition. □

Recall the bounded cancellation constant $\mathrm{BCC}(f)$. Suppose $\alpha \cdot \beta$ is the concatenation

of two legal paths. Then we have $\mathrm{len}([f(\alpha \cdot \beta)]) > \lambda \,\mathrm{len}(\alpha \cdot \beta) - 2\,\mathrm{BCC}(f)$.

We can deduce a bound on the length of Nielsen paths. Let $C_f := \frac{2\,\mathrm{BCC}(f)}{\lambda - 1}$. If α, β are legal and $\mathrm{len}(\alpha \cdot \beta) > C_f$ then $\mathrm{len}([f(\alpha \cdot \beta)]) > \mathrm{len}(\alpha \cdot \beta)$ so the length of $([f^n(\alpha \cdot \beta)])$ is strictly increasing. Therefore, if $\alpha \cdot \beta$ is an indivisible Nielsen path, then $\mathrm{len}(\alpha) \leq C_f, \mathrm{len}(\beta) \leq C_f$.

The following fact is a well-known fact for the free group case and also applies here, mainly because trees are also locally finite in the GBS case.

Lemma 2.3.3. *There are finitely many orbits of periodic indivisible Nielsen paths.*

Remark 2.3.4. Since T is locally finite, there are finitely many orbits of edge paths with bounded length. However in general the endpoints of pINPs are not vertices of T so pINPs are not edge paths.

Before proving Lemma 2.3.3 we introduce the following notion which will enable us to work with edge paths while looking for periodic indivisible Nielsen paths.

Definition 2.3.5. A *pseudo-pINP* is an edge path $\gamma \subset T$ with one illegal turn such that there exists $n \geq 1$ and $g \in G$ such that $\gamma \subset g[f^n(\gamma)]$.

Lemma 2.3.6. *Every pseudo-pINP contains a unique pINP. Conversely, the minimal edge path containing a pINP is a pseudo-pINP.*

Proof. Suppose γ is an edge path with one illegal turn, and let $n \geq 1$ and $g \in G$ be such that $\gamma \subset g[f^n(\gamma)]$. Up to replacing f by gf^n we may suppose $\gamma \subset [f(\gamma)]$. Write γ as the concatenation $\alpha \cdot \eta_1 \cdot \eta_2 \cdot \beta$ where the illegal turn is between η_1 and η_2, and η_1 and η_2 are the maximal subpaths such that $f(\eta_1) = f(\bar{\eta}_2)$. If γ contains a pINP γ' then the length of each legal branch of γ' must be $\frac{\lambda \,\mathrm{len}(\eta_1)}{\lambda - 1}$, which proves uniqueness. For the existence, since $\gamma \subset f(\gamma)$ this length defines a unique subpath $\gamma' \subset \gamma$ and $[f(\gamma')] = \gamma'$.

Conversely, the minimal edge path γ containing a pINP γ' is sent to an edge path containing $gf^n(\gamma')$ for some $n \geq 1$ and $g \in G$. Since $\gamma' \subset g[f^n(\gamma')] \subset g[f^n(\gamma)]$ and γ is the minimal edge path containing γ', then γ is a pseudo-pINP. □

Now we prove Lemma 2.3.3.

Proof of Lemma 2.3.3. First of all the tree T is locally finite and there are finitely many orbits of illegal turns. There is a bound on the lengths of pINPs, namely the critical constant C_f. Finally every pINP crosses a unique illegal turn.

Every pINP is contained in a pseudo-pINP whose legal branches have length at most $C_f + \max_{e \in E(T)} \mathrm{len}(e)$. There are finitely many orbits of edge paths with bounded length, hence finitely many orbits of pseudo-pINPs with bounded length. Since every pseudo-pINP contains a unique pINP by Lemma 2.3.6, there are finitely many orbits of pINPs. □

The important implication of Lemma 2.3.3 is that the set of orbits of indivisible Nielsen paths can be computed algorithmically. To prove this for GBS groups we use an approach which resembles [CL15, Sections 5, 6], where *long turns* are used to understand pINPs.

The algorithm which finds all periodic indivisible Nielsen paths relies on the correspondance with pseudo-pINPs given by Lemma 2.3.6. It suffices to prove that all pseudo-pINPs can be computed algorithmically. This would be quite simple if we knew the bound C_f. However we do not know how to estimate $\mathrm{BCC}(f)$ and instead we will give an algorithm which does not need it. The algorithm consists in testing all edge paths with bounded length for pseudo-pINPs, and using a criterion to make sure that the bound on the length was sufficiently big. If not then the bound is increased: eventually it becomes greater than C_f and the criterion is satisfied.

Let $\gamma = \alpha \cdot \beta$ be an edge path with $\mathrm{len}(\alpha), \mathrm{len}(\beta) \neq 0$ such that α, β are legal subpaths and the turn at the concatenation is illegal.

The following technical lemma studies the behaviour of the sequences $([f^n(\gamma)])_{n \in \mathbb{N}}$.

Lemma 2.3.7. *(i) There are four mutually exclusive behaviours for the sequence $([f^n(\gamma)])$, illustrated by Figure 2.3:*

- *(1) for some $n \in \mathbb{N}$, $[f^n(\gamma)]$ is legal and neither $f^n(\alpha)$ nor $f^n(\beta)$ contains the other*
- *(2) for some $n \in \mathbb{N}$, $f^n(\alpha) \subset f^n(\beta)$ or $f^n(\beta) \subset f^n(\alpha)$*
- *(3) γ contains a pseudo-pINP*
- *(4) γ does not contain any pseudo-pINP but there exists $k \geq 1$ such that $[f^k(\gamma)]$ does*

(ii) There is an algorithm which takes an edge path γ with one illegal turn and returns what case $([f^n(\gamma)])$ belongs to

(iii) If α, β are both longer than C_f then case (2) is not possible

(iv) If there exists a pINP η such that the minimal edge path containing η strictly contains γ then γ is in case (2).

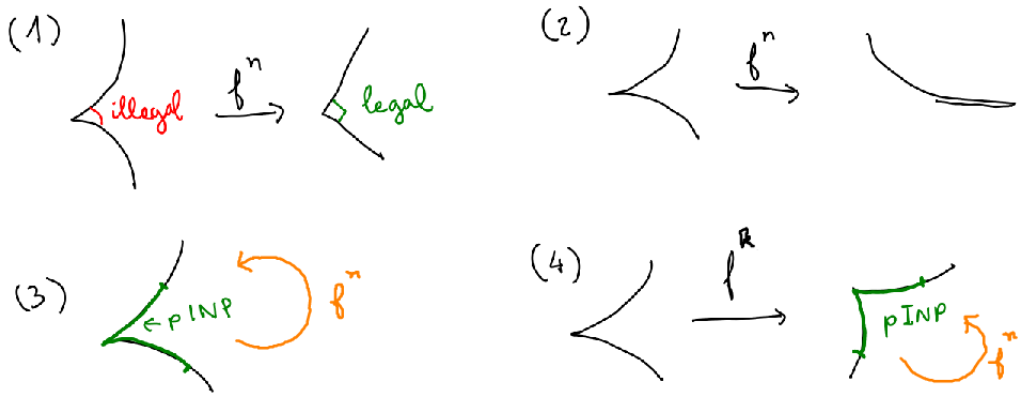

Figure 2.3 – Four different cases for the behaviour of $[f^n(\gamma)]$

Proof. (i) If there exists $n \in \mathbb{N}$ such that $[f^n(\gamma)]$ is legal then the sequence belongs to case (1) or (2). Otherwise, for all $n \in \mathbb{N}$, the path $[f^n(\gamma)]$ contains an illegal turn. In particular its length is never zero. Let $B \geq C_f$ be a bound on the length of legal branches of pseudo-pINPs. There are finitely many edge paths with length at most $2B$ so there exist $k \geq 0, n \geq 1$ and $g \in G$ such that $[f^k(\gamma)]$ and $g[f^{k+n}(\gamma)]$ are either equal, or intersect on a length at least B on each side of the illegal turn. Let η be the maximal edge path in the intersection $[f^k(\gamma)] \cap g[f^{k+n}(\gamma)]$. Let α_η be one of the legal branches of η. Since $\mathrm{len}(\alpha_\eta) \geq C_f$ we have $\mathrm{len}(f^n(\alpha_\eta) \cap [f^n(\eta)]) \geq \mathrm{len}(\alpha_\eta)$. Therefore the image $[f^n(\eta)]$ contains $g^{-1}\eta$, so η is a pseudo-pINP. The sequence belongs to case (3) if $k = 0$ and to case (4) otherwise.

Case (1) obviously excludes cases (2), (3) and (4). Case (2) excludes cases (3) and (4) since it implies that $[f^n(\alpha \cdot \beta)]$ is a legal path. Definition of cases (3) and (4) also implies that they are mutually exclusive.

(ii) Before giving the algorithm which determines the case, let us give some preliminaries:

— By Proposition 2.1.4 (v), there is an algorithm which, given finite non-backtracking edge paths α, β, decides whether they belong to the same orbit or whether there exists $g \in G$ such that $\alpha \subset g\beta$.

— There is an algorithm which takes an edge path η and a period $p \geq 1$, and determines if η is a pseudo-pINP with period at most p. Compute the iterates $[f^n(\eta)]$ for $n \in \{1, \ldots, p\}$. If there exists n such that $[f^n(\eta)]$ contains a translate of η then η is a pseudo-pINP whose period divides n. Otherwise it is not a pseudo-pINP with period at most p.

Let γ be an edge path with one illegal turn. Repeat the following steps, starting with $n = 0$:

— Compute $f^n(\gamma)$.

— If $f^n(\gamma)$ is legal, check whether $f^n(\alpha) \subset f^n(\beta)$ or vice-versa. If yes then γ belongs to case (2), else it belongs to case (1). The algorithm stops.

— If $f^n(\gamma)$ is not legal, check if it contains a pseudo-pINP with period at most $n+1$. Suppose it contains a pseudo-pINP with period $p \leq n+1$. Then either case (3) or case (4) is true. The pINP contained in η also has period p: if γ contains a pINP η_0 then η contains its f^n-image, so η_0 has period p. Then it suffices to check if γ contains any pseudo-pINP with period p. If yes then γ is in case (3), else it is in case (4).

— Replace n by $n+1$.

Because of (i) this will eventually terminate.

(iii) If $\mathrm{len}(\alpha) \geq C_f$ then for every $n \in \mathbb{N}$, $\mathrm{len}(f^n(\alpha) \cap [f^n(\gamma)]) > \lambda\,\mathrm{len}(\alpha) - \mathrm{BCC}(f) > \mathrm{len}(\alpha)$. A similar statement holds if $\mathrm{len}(\beta) \geq C_f$. When both α and β are longer than C_f, (2) cannot occur.

(iv) Suppose that there exists a pINP η, a pseudo-pINP η' containing η, such that $\gamma \subsetneq \eta'$, and that γ does not contain η. One of the endpoints x of γ lies in the interior of η. Let y be the endpoint of η in the same legal branch as x. The distance $d_T(f^n(x), f^n(y))$ increases exponentially since $[x, y]$ is legal. The distance between $f^n(y)$ and the illegal turn is bounded by C_f. For n_0 big enough, $d_T(f^{n_0}(x), f^{n_0}(y)) > C_f$ so $f^{n_0}(x)$ lies in the simplified part of $f^{n_0}(\eta)$. Since $\gamma \subset \eta'$, $[f^{n_0}(\gamma)]$ is contained in $[f^{n_0}(\alpha)]$ or $[f^{n_0}(\beta)]$ so γ is in case (2). □

Proposition 2.3.8. *There is an algorithm taking a train track map $f : T \to T$ and finding all orbits of minimal pseudo-pINPs.*

Proof. Proposition 2.1.4 (v) implies that given two edge paths, one can decide whether they belong to the same orbit. As a result one can list all orbits of edge paths of given length: it suffices to choose a representative for each orbit of vertex, and then construct all edge paths with given length starting with these vertices.

Start with $L = 1$. Apply the following steps.

— Let $\mathcal{L}_L$ be a list of representatives of edge paths with shape $\alpha \cdot \beta$ where α and β are legal paths with combinatorial length at most L.

— For each path in $\mathcal{L}_L$, determine what case they belong to using Lemma 2.3.7 (ii).

— If there exist a path in case (2) then increase L by 1 and start again. Else stop. Let $\mathcal{T}$ be the subset of $\mathcal{L}_L$ consisting of paths in case (3).

Every path $\gamma \in \mathcal{T}$ contains a pseudo-pINP, moreover point (iii) of Lemma 2.3.7 ensures that we can find such a pseudo-pINP $\eta_0 \subset \gamma$ and its period p. For every other pseudo-pINP $\eta \subset \gamma$, the period of η is also p. Compute $[f^p(\eta)]$ for every $\eta \subset \gamma$ in order to find all pseudo-pINPs in γ and find the minimal one.

This algorithm eventually stops because after enough steps L becomes greater than C_f so case (2) does not occur. Suppose that the algorithm stops, then for every minimal pseudo-pINP η there exists $\gamma \in \mathcal{T}$ such that $\eta \subset \gamma$. By contradiction suppose otherwise. Then $\text{len}(\eta) > L$. There exists a subpath $\gamma \subsetneq \eta$ of combinatorial length L and by minimality of η, γ does not contain the pINP contained in η. Then by (iv) of Lemma 2.3.7 γ is in case (2) which is a contradiction. □

Of course it is more convenient to work with edge paths than with arbitrary paths which can start or end in the middle of an edge, especially for algorithmic purposes. Instead of using pseudo-pINPs, we will subdivide T at the endpoints of pINPs so that all pINPs become actual edge paths:

Lemma 2.3.9. *There exists a subdivision $s : T \to T'$ with $T' \in \mathcal{D}$ and a train track map $f' : T' \to T'$ such that $f' \circ s = s \circ f$, such that the endpoints of all pINPs for f' are vertices of T'.*

Proof. Define T' by subdividing T at every endpoint of periodic indivisible Nielsen paths. Since there are finitely many orbits of pINPs the tree T' is simplicial. Let $s : T \to T'$ be the corresponding isometry.

Since the set of pINPs is stable under f, the map f' induced by f on T' maps vertex to vertex. The map f' is a train track map. □

Lemma 2.3.10. *One can compute a subdivision $s : T \to T'$ and a train track map $f' : T' \to T'$ such that $f' \circ s = s \circ f$ and such that all pINPs for f' start and end at vertices of T'.*

Proof. The map f exists by Lemma 2.3.9. We will give an algorithmic construction.

One can find all orbits of minimal pseudo-pINPs by Proposition 2.3.8. Let γ be a minimal pseudo-pINP. It contains a unique pINP η. There exists $n \geq 1$ and $g \in G$ such

that $\eta = g[f^n(\eta)]$. The first point of η (resp. last point) is a vertex if and only if the first edge (resp. last edge) of γ is equal to the first edge (resp. last edge) of $g[f^n(\gamma)]$. This can be checked algorithmically.

There exists $k \geq 1$ such that all pINPs are k-periodic. Step by step, we will construct a subdivision of T which is f^k-invariant. We will check that it is also f-invariant.

Start with $S = T$.

While there exists a pINP $\eta = [x, y]$ for f^k such that x is not a vertex, define the subdivided tree S' as follows. Let e be the edge of S which contains x. Subdivide e by adding a new vertex v representing x. The map f^k induces a map on S' with $f^k(x) = gx$ where g is such that $f^k(\eta) = g\eta$.

By repeating this with $S := S'$ for every pINP, we obtain an f^k-invariant subdivision T'.

Now we would like to prove that the subdivision is f-invariant. We need to define f on vertices of $V(T') \setminus V(T)$. Suppose $v \in V(T') \setminus V(T)$. In T the point v is not a vertex but it is the endpoint of a pINP η. There exists a unique pINP η' such that $[f(\eta)] = \eta'$, and the path η' can be computed. The map f sends the endpoints of η to those of η': define $f'(v)$ as the corresponding endpoint of η' in T'. By construction $f'(v)$ is a vertex.

The construction does not depend on the choice of η since it exhibits the subdivision of Lemma 2.3.9. □

2.3.2 Pseudo-periodic conjugacy classes of G

Recall that the conjugacy class of $g \in G$ is *pseudo-periodic* for $\Phi \in \mathrm{Out}(G)$ if $\|\phi^n(g)\|_T$ is bounded for $n \to \infty$ and $\phi \in \Phi$.

The motivation for the notion of pseudo-periodicity is given by the following result, which we will prove below:

Proposition 2.3.11. *If there exists a simple element $h \in G$ such that h is pseudo-periodic, then a power of Φ is reducible.*

Lemma 2.3.12. *Let $g \in G$ be a loxodromic element. Let $\phi \in \mathrm{Aut}(G)$: The following conditions are equivalent:*

— *$\|\phi^n(g)\|_T$ does not tend to infinity*

— *$\|\phi^n(g)\|_T$ is bounded.*

Proof. The axis of g in T can be partitioned into maximal legal segments, concatenated at illegal turns. Let $s_1, \dots, s_m$ be consecutive maximal legal segments forming a fundamental domain of $\mathrm{Axis}_T(g)$.

For every $n \in \mathbb{N}$ there exists $I_n \subset \{1, \dots, m\}$ such that $\bigcup_{i \in I_n} f^n(s_i)$ contains a fundamental domain of the axis of $\phi^n(g)$. We may assume $I_n \subset I_{n-1}$ for every $n \in \mathbb{N}$. For $i \in I_n$ let $J_i^n := f^n(s_i) \cap \mathrm{Axis}_T(\phi^n(g))$. It is a legal segment.

If $\|\phi^n(g)\|_T$ is not bounded then at least one of the J_i^n has unbounded length. Suppose for $N \in \mathbb{N}$, $\mathrm{len}(J_{i_0}^N) > 2C_f$. When applying f to $J_{i_0}^n$, it is stretched by a factor λ, however there might be cancellation at the ends of $J_{i_0}^n$ because of illegal turns. Since this cancellation cannot exceed $\mathrm{BCC}(f)$, we have $\mathrm{len}(J_{i_0}^{N+1}) \geq \lambda \,\mathrm{len}(J_{i_0}^N) - 2\,\mathrm{BCC}(f) > \frac{\lambda}{2}\,\mathrm{len}(J_{i_0}^N) + \lambda C_f - 2\,\mathrm{BCC}(f) > \frac{\lambda}{2}\,\mathrm{len}(J_{i_0}^N)$ so $\mathrm{len}(J_{i_0}^n) \geq \left(\frac{\lambda}{2}\right)^{n-N} \mathrm{len}(J_{i_0}^N)$. Therefore $\|\phi^n(g)\|_T$ goes to infinity. □

Remark 2.3.13. Equivalently the conjugacy class of an element $g \in G$ is pseudo-periodic for Φ if there exists $S \in \mathcal{D}$ such that $\|\phi^n(g)\|_S$ is bounded with $\phi \in \Phi$, and equivalently if for all $S \in G$, $\|\phi^n(g)\|_S$ is bounded.

Pseudo-periodic elements are the right analogue of periodic elements in the GBS context. Although the conjugacy class of a periodic element is not periodic, the translation length is.

Lemma 2.3.14. *Let $\phi \in \mathrm{Aut}(G)$. Let $h, g \in G$ be loxodromic elements. Suppose there exists $t \in \mathcal{D}$ and $v \in \mathrm{Axis}_T(h) \cap \mathrm{Axis}_T(g)$ such that $[v, h^2 v] = [v, g^2 v]$. Then in any $S \in \mathcal{D}$ there exists $w \in S$ such that $[w, g^2 w] = [w, h^2 w]$ is a common pair of fundamental domains for g and h. In particular $\|g\|_S = \|h\|_S$ for all $S \in \mathcal{D}$.*

Proof. Let ϕ, h, g, T, v be as described in the hypotheses.

Let $f : T \to S$ be a G-equivariant application. See Figure 2.4. The segment $[f(v), hf(v)]$ intersects $\mathrm{Axis}_S(h)$. Let l_h be the distance between $f(v)$ and $\mathrm{Axis}_S(h)$. It is equal to the distance between $hf(v)$ and the axis. Moreover the length of the intersection $[f(v), hf(v)] \cap [hf(v), h^2 f(v)]$ is equal to l_h. The same goes for l_g and since $gv = hv$ and $g^2 v = h^2 v$, we have $l_g = l_h$.

Let $w_S \in S$ be the point of $[f(v), hf(v)]$ at distance $l_h = l_g$ from $f(v)$. It belongs to $\mathrm{Axis}_S(h)$ and $\mathrm{Axis}_S(g)$.

The translation length of both g and h in S is $d_S(f(v), hf(v)) - 2l_h$. Therefore $hw_S = gw_S$ is the point of $[f(v), hf(v)]$ at distance l_h from $hf(v)$ and $h^2 w_S = g^2 w_S$ is the point at distance l_h from $h^2(v)$. □

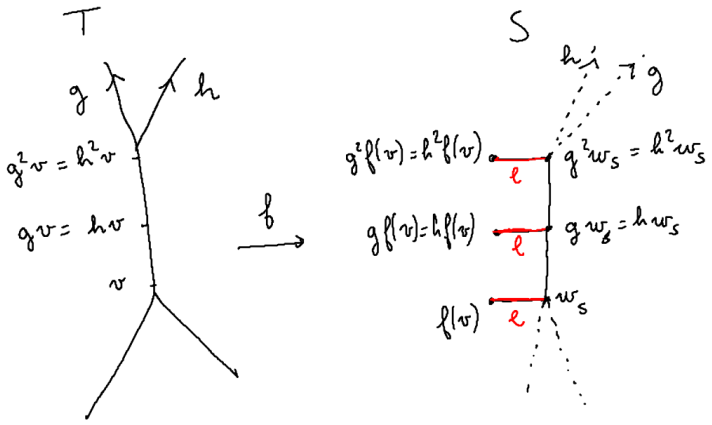

Figure 2.4 – Description of the proof of Lemma 2.3.14: image of the pair of fundamental domains in the new tree S

Remark 2.3.15. The previous lemma is also true when replacing 2 with any $n \geq 2$: if g, h share n consecutive fundamental domain in some $T \in \mathcal{D}$ then they do in every $T \in \mathcal{D}$.

Corollary 2.3.16. *Let $h, g \in G$ be loxodromic elements such that there exists $T \in \mathcal{D}$ and $v \in \mathrm{Axis}_T(h) \cap \mathrm{Axis}_T(g)$ such that $[v, h^2 v] = [v, g^2 v]$. Then any special factor containing h also contains g.*

Proof. Let A be a special factor containing h. Let $S \in \mathcal{D}^{\mathcal{A}}$ be a tree such that a collapse $\bar{S}$ of S has a vertex x with stabilizer A. By Lemma 2.3.14 the axes of h and g in S share a fundamental domain $[w, gw] \subset S$.

In $\bar{S}$, A is the stabilizer of the vertex x. Therefore the axis of h is sent to x by the collapse map $S \to \bar{S}$. The fundamental domain of g in the axis of h is also sent to x, so the whole axis of g is sent to x by equivariance. Thus g fixes x and $g \in A$. □

In the following result, the fact that trees in $\mathcal{D}$ are locally finite is fundamental.

Proposition 2.3.17. *Let $\phi \in \mathrm{Aut}(G)$. Let $g \in G$ be a loxodromic element. The following conditions are equivalent:*

(i) the conjugacy class of g is pseudo-periodic for the outer class $\Phi = [\phi]$

(ii) for any $T \in \mathcal{D}$, for all $N \geq 1$, there exist distinct $n, m \in \mathbb{N}$ such that $\phi^n(g)$ and $\phi^m(g)$ admit N common consecutive fundamental domains up to translation

(iii) for all $N \geq 1$, there exists $m > 0$ such that for any $S \in \mathcal{D}$, g and $\phi^m(g)$ admit N common consecutive fundamental domains up to translation in S

(iv) for any $T \in \mathcal{D}$ the sequence $(\|\phi^n(g)\|_T)_{n\to-\infty}$ is bounded.

Proof. Let us prove (i) $\Rightarrow$ (ii). Suppose the conjugacy class of g is pseudo-periodic.

Let $f : T \to T$ be a representative for ϕ. Let $B > \sup_{n\in\mathbb{N}} \|\phi^n(g)\|_T$. Let $N \in \mathbb{N}$. Since there are finitely many orbits of edge paths with length at most NB, there exist distinct $n, m \in \mathbb{N}$ and $h \in G$ such that $\phi^n(g)$ and $h\phi^m(g)h^{-1}$ share N consecutive fundamental domains.

To prove (ii) $\Rightarrow$ (iii), take $T \in \mathcal{D}$. By (ii) there exists n, m such that $\phi^n(g)$ and $\phi^m(g)$ have N consecutive fundamental domains up to translation. By Lemma 2.3.14 and Remark 2.3.15, these two elements have N consecutive fundamental domains, up to translation, in every tree in $S \in \mathcal{D}$. In particular in $S \cdot \phi^{-n}$ they do. This implies that g and $\phi^{m-n}(g)$ share N fundamental domains up to translation in S, for every $S \in \mathcal{D}$.

Suppose (iii): in particular, by taking $S' := S \cdot \phi^n$ we have $\|\phi^n(g)\|_S = \|\phi^{n+m}(g)\|_S$ for any $n \in \mathbb{Z}$. Therefore $\|\phi^n(g)\|_S$ is bounded so g is pseudo-periodic for $[\phi]$ so (i) and (iv) are true.

Finally suppose (iv). Condition (iv) is condition (i) for ϕ^{-1} instead of ϕ, and it implies (iv) for ϕ^{-1}, which is (i) for ϕ. Therefore (iv) $\Rightarrow$ (i). □

Now we can prove Proposition 2.3.11, which states that the minimal factor containing an element whose conjugacy class is pseudo-periodic is itself periodic.

Proof of Proposition 2.3.11. Condition (iii) of Proposition 2.3.17 states that there exist n, m such that h and $\phi^m(h)$ share two consecutive fundamental domains up to translation. Using Corollary 2.3.16 this implies that any special factor containing h also contains a conjugate of $\phi^m(h)$, and vice versa. Thus the minimal factors containing these elements are conjugate and the conjugacy class of these factors is ϕ^m-periodic. □

The following results are the key for finding pseudo-periodic elements in G:

Proposition 2.3.18. *Let $g \in G$ be an element whose conjugacy class is pseudo-periodic for $\Phi \in \mathrm{Out}(G)$. Suppose $f : T \to T$ is a train-track representative for Φ. Then the axis of g in T is a concatenation of periodic indivisible Nielsen paths.*

Proof. Let $\phi \in \Phi$ be such that f is ϕ-equivariant. The axis of g in T is a concatenation of maximal legal subsegments interrupted by illegal turns. The tightened image of the axis of g by f^n is the axis of $\phi^n(g)$. Since f^n maps legal segments to legal segments, the number N_n of orbits of maximal legal segments in $\mathrm{Axis}_T(\phi^n(g))$ never increases with n.

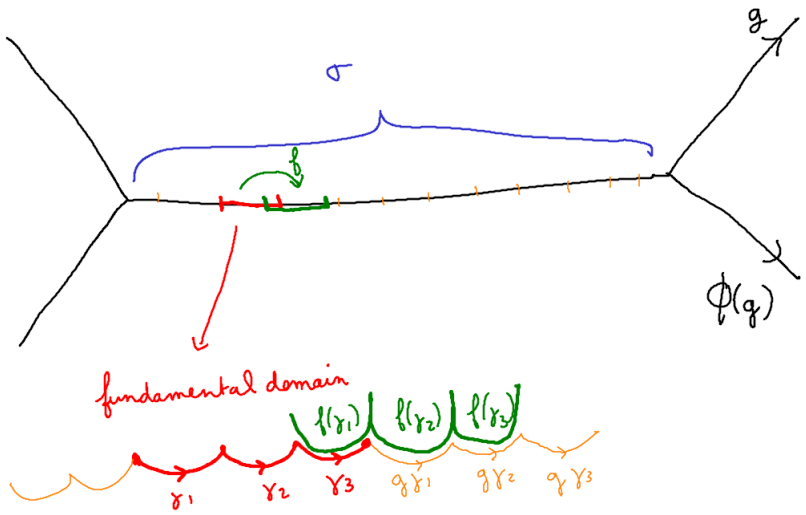

Figure 2.5 – Picture of the action of f on legal segments

Since g and $\phi^{nk}(g)$ share two fundamental domains up to translation (Proposition 2.3.17 (iii)) for some fixed k and every $n \in \mathbb{N}$, the number N_n is actually a constant N.

By Proposition 2.3.17 (iii) there exists $n \geq 1$ and $h \in G$ such that $h\phi^n(g)h^{-1}$ and g have at least $N+1$ consecutive fundamental domains in common. Up to replacing ϕ by $c_h \circ \phi^n$ for some $l \in \mathbb{Z}$, and f by hf^n, we may suppose that g and $\phi(g)$ share $N+1$ consecutive fundamental domains. Let us call σ the segment where both axes overlap. Up to replacing ϕ by $c_{\phi(g^l)} \circ \phi$ we may suppose that for any fundamental domain $\eta \subset \sigma$, the first point of $f(\eta) \cap \sigma$ is contained in η (see Figure 2.5). Note that the set of pINPs for ϕ does not change when replacing ϕ with a power or composing with an inner automorphism, so these assumptions will not change the outcome of the proof.

Let $\gamma_1, \dots, \gamma_N, \dots, \gamma_{N^2}, \gamma_{N^2+1}$ be consecutive maximal legal segments which appear whole in σ. The isometry g shifts the legal subsegments: for $i \in \{1, \dots, N^2+1-N\}$ we have $\gamma_{i+N} = g\gamma_i = \tilde{\phi}(g)\gamma_i$.

The map f sends legal paths to legal paths, and it induces a bijection between the sets of maximal legal subsegments of $\mathrm{Axis}_T(g)$ and $\mathrm{Axis}_T(\phi(g))$. This bijection preserves the order of the segments. On the part where the axes overlap, f shifts the subsegments γ_i by an amount which does not depend on i. More accurately, there exists $0 \leq j < N$ such that for every $1 \leq i \leq N(N-1)$, $f(\gamma_i) \supset \gamma_{i+j}$.

In particular, for $i \in \{1, \dots, N+1\}$ we have

$$f^N(\gamma_i) \supset f^{N-1}(\gamma_{i+j}) \supset \cdots \supset \gamma_{i+Nj} = g^j\gamma_i.$$

Since for all $i \in \{1, \dots, N+1\}$ we have $g^j \gamma_i \subset f^N(\gamma_i)$ and f^N stretches legal segments uniformly, there exists a unique point $p_i \in \gamma_i$ such that $f^N(p_i) = g^j p_i$. Besides $p_{N+1} = g p_1$. Thus for $1 \leq i \leq N$ the segment $[p_i, p_{i+1}]$ is a periodic Nielsen path, and since it has a unique illegal turn it is a pINP for f. Therefore $\mathrm{Axis}_T(g)$ is a concatenation of pINPs. □

We also have a converse:

Proposition 2.3.19. *Suppose that $g \in G$ is a loxodromic element whose axis is a concatenation of pINPs. Then g is pseudo-periodic.*

Proof. Suppose the axis of $g \in G$ is a concatenation of pINPS. There exists a fundamental $\eta_1 \cdot \dots \cdot \eta_k$ for g where η_i is a pINP for $i \in \{1, \dots, k\}$. There exists a common $n \in \mathbb{N}$ such that for all $i \in \{1, \dots, k\}$ there exists $h_i \in G$ such that $[f^n(\eta_i)] = h_i \eta_i$. By continuity of f^n the paths $h_i \eta_i$ and $h_{i+1} \eta_i$ have a common endpoint, although h_i may be different from h_{i+1}.

Therefore, for every $k \in \mathbb{N}$, $\mathrm{len}([f^{nk}(\eta_1 \cdot \dots \cdot \eta_k)]] \leq \mathrm{len}(\eta_1 \cdot \dots \cdot \eta_k)$. Since f is Lipschitz, this proves that $\|\phi^m(g)\|_T \to \infty$ when $n \to \infty$. □

Remark 2.3.20. A consequence is that there cannot be any illegal concatenation of pINPs forming the axis of a loxodromic element of G. If there were such an element g, up to taking a power of Φ we could suppose the axis of g is made of INPs with an illegal concatenation occuring at some point. Then the sequence $\|\phi^n(g)\|_T$ would be non increasing and non constant because of the illegal turn. This contradicts Proposition 2.3.17 (iv).

2.4 Pseudo-periodic subgroups

In this section we introduce a collection of subgroups whose loxodromic elements are exactly the elements of G whose conjugacy class is pseudo-periodic, and we give an algorithm which computes these subgroups.

These groups can be understood as *pseudo-periodic subgroups*, which are analogous to fixed subgroups for free groups. The notion of pseudo-periodicity which is used here depends on a choice of actual automorphism $\phi \in Aut(G)$ and is not a conjugacy class invariant.

2.4.1 Algorithmic computation of pseudo-periodic conjugacy classes

Fix $\Phi \in \mathrm{Out}(G)$. Our aim in this subsection is to determine a subset of points which can be effectively constructed, and such that the loxodromic elements of its stabilizer are pseudo-periodic.

Let $f : T \to T$ be a train track map for Φ. Let $\eta := [x_0, x_1] \subset T$ be a periodic indivisible Nielsen path for f. The *Nielsen class* associated to η is the set

$$VY(\eta) := \{y \in T / \exists x_0 = y_0, y_1, \dots, y_n \text{ s.t. } \forall 1 \leq i \leq n \quad [y_{i-1}, y_i] \text{ is a pINP } \}$$

The set of pINPs with both ends in $VY(\eta)$ is called $EY(\eta)$. For any $\eta' \in EY(\eta)$ we have $VY(\eta') = VY(\eta)$.

Lemma 2.4.1. *Let $\Phi \in \mathrm{Out}(G)$ and let $f : T \to T$ be a representative for Φ. Let η be a pINP for f. Then for $g \in G$ the following are equivalent:*

(i) $VY(\eta) = gVY(\eta)$

(ii) $VY(\eta) \cap gVY(\eta) \neq \varnothing$

Proof. The implication (i) $\Rightarrow$ (ii) is obvious since $VY(\eta) \neq \varnothing$.

Suppose (ii). Let us prove (i). Let $x \in VY(\eta) \cap gVY(\eta)$. Let $y \in VY(\eta)$. By definition of $VY(\eta)$, there is a concatenation of pINPs from x to y. Since $x \in gVY(\eta)$ there is also a concatenation from x to gy. Thus y and gy can be connected by a concatenation of pINP so they belong to the same Nielsen class. We deduce $gVY(\eta) = VY(\eta)$. □

Lemma 2.4.2. *Let $\Phi \in \mathrm{Out}(G)$ with a train track representative $f : T \to T$. Let η be a pINP for Φ.*

Suppose $\mathrm{Stab}(VY(\eta))$ *contains a loxodromic element.*

The action of $\mathrm{Stab}(VY(\eta))$ *on the subtree* $Y(\eta) := \mathrm{conv}(VY(\eta))$ *is cocompact.*

Proof. Let $VY := VY(\eta)$, $Y := \mathrm{conv}(VY)$ and $EY := EY(\eta)$.

The set $VY/\mathrm{Stab}(VY)$ is finite because the image of VY in T/G is finite: by Lemma 2.4.1, if $g \in G$, $y \in VY$ and $gy \in VY$, then $g \in \mathrm{Stab}(VY)$.

The set $EY/\mathrm{Stab}(VY)$ is finite because EY/G is finite and for $g \in G$, $VY \cap gVY \neq \varnothing \Rightarrow g \in \mathrm{Stab}(VY)$. There exist representatives $I_1, \dots, I_k$ for EY and their union $I_1 \cup \cdots \cup I_k$ is a compact subset of T whose orbit covers Y.

Thus $Y/\mathrm{Stab}(VY)$ is compact. □

Corollary 2.4.3. *Let $\Phi \in \mathrm{Out}(G)$ with a train track representative $f : T \to T$ be such that there exists a pINP η for f. The sets $VY(\eta)/\mathrm{Stab}(VY(\eta))$ and $EY(\eta)/\mathrm{Stab}(VY(\eta))$ are finite.*

The stabilizers of the Nielsen classes contain all the information on pseudo-periodic conjugacy classes, and they can be computed:

Proposition 2.4.4. *Let $\Phi \in \mathrm{Out}(G)$. Let $f : T \to T$ be a train track representative for Φ. For every pINP η for f, one can compute algorithmically a description of the set $VY(\eta)$ and of its stabilizer in the following form:*

- *a finite set of generators for the subgroup $\mathrm{Stab}(VY(\eta))$*
- *a list of representatives for each $\mathrm{Stab}(VY(\eta))$-orbit of points of $VY(\eta)$*
- *a list of representatives for each $\mathrm{Stab}(VY(\eta))$-orbit of pINP in $EY(\eta)$.*

Proof. By Lemma 2.3.10 we may subdivide T such that the endpoints of all pINPs are vertices.

Let η be a pINP for f.

We construct the set $EY := EY(\eta)$ as follows. It suffices to construct all possible pINP concatenations starting from an endpoint of η, which actually constructs the vertices of $VY := VY(\eta)$.

Define EY_n as the list of pINPs which appear in concatenations of length at most n from y_0, where $\eta = [y_0, y_1]$. For all $n \geq 1$ we have $EY_n \subset EY_{n+1}$.

The construction of EY_n is algorithmic. One needs a list of representatives for all G-orbits of pINPs, which can be computed using Proposition 2.3.8. Given a pINP γ ending with a vertex x, we need to know whether there is a pINP $\gamma' \neq \gamma$ with x as endpoint. This question can be answered using these facts, which are consequences of Proposition 2.1.4:

- given a G-orbit of pINPs $[\gamma']$ one can decide whether one endpoint of $[\gamma']$ is in the same orbit as x
- if so, one can compute a representative γ' with x as an endpoint and compute all other translates by applying repeatedly a generator of G_x. Since T is locally finite, $G_x \cdot \gamma'$ is finite so eventually, all translates of γ' starting at x have been computed.

Since there are finitely many G-orbits of pINPs there exists $n \in \mathbb{N}$ such that all orbits of pINPs in EY_{n+1} also appear in EY_n.

Choose a minimal set of representatives in $\mathcal{R} \subset EY_n$ such that all orbits of pINPs in EY_{n+1} are represented in $\mathcal{R}$. Let $\eta_1, \dots, \eta_s$ be the pINPs of $EY_n \setminus \mathcal{R}$. For every $1 \leq i \leq s$ there exists an element $g_i \in G$ such that $g_i\eta_i$ belongs to $\mathcal{R} \subset EY_{n-1}$. Let a be a generator of $\mathrm{Stab}(\eta)$. Let $G_\eta := \langle g_1, \dots, g_s, a\rangle$.

Then define $EY_\eta := G_\eta \cdot EY_n$.

Let us prove that $EY_\eta = EY$. It suffices to prove that EY_η is the set of all pINPs which belong to a pINP path to y_0.

Suppose $\gamma \in EY_\eta$. There exists a word $g_1 \dots g_m$ in the generators of G_η and a pINP $\gamma_0 \in EY_n$ such that $\gamma = g_1 \dots g_m\gamma_0$.

For every $1 \leq i \leq m$, the paths $\gamma_0, g_i\gamma_0$ belong to EY_n and can be joined by a concatenation of pINPs within EY_n. Thus there is a concatenation containing all pINPs

$$\gamma_0, g_1\gamma_0, g_1g_2\gamma_0, \dots, g_1 \dots g_m\gamma_0 = \gamma.$$

Conversely suppose that there exists a concatenation of pINPs $\gamma_0, \gamma_1, \dots, \gamma_m$ such that $\gamma_0 \in EY_n$. Let us prove that $\gamma_m \in EY_\eta$. We proceed by induction on m. The case $m = 0$ is obvious. For greater m suppose every concatenation of at most m pINPs starting with a pINP in EY_n is contained in EY_η. By the induction hypothesis, there exists $g \in G_\eta$ such that $g\gamma_{m-1}$ belongs to EY_n. Thus $g\gamma_m$ belongs to EY_{n+1} so there exists $i \in \{1, \dots, s\}$ such that $g_ig\gamma_m \in EY_n$. Since $g_ig \in G_\eta$, this proves that $\gamma_m \in EY_\eta$. Thus a concatenation of $m + 1$ pINPs belongs to EY_η. We can conclude by induction.

Now let us prove that the stabilizer of EY_η is the subgroup G_η, which is finitely generated. The inclusion $G_\eta \subset \mathrm{Stab}(EY_\eta)$ is true. Conversely let $g \in \mathrm{Stab}(EY_\eta)$. There exists a concatenation of pINPs $\eta, \dots, g \cdot \eta$. There exists $g' \in G_\eta$ such that $g'\eta = g \cdot \eta$ so $g \in g' \,\mathrm{Stab}(\eta) \subset G_\eta$.

The list EY_n provides the list of representatives of the orbits of pINPs in EY. Moreover the ends of the elements of EY_n are a list VY_n of representatives of the orbits of VY. There may be redundant elements in both EY_n and VY_n and they may be eliminated algorithmically using Proposition 2.1.4. □

Serre's Lemma ([Ser77, Section 6.5, Corollary 2]) gives a criterion for deciding whether a finitely generated subgroup acting on a tree is elliptic.

Lemma 2.4.5. *Suppose a finitely generated group G acts by isometries on a simplicial tree T. Let $\{s_1, \dots, s_n\}$ be a finite generating set for G. Then the action of G is elliptic*

if and only if

— *for all* $i \in \{1, \dots, n\}$ *the isometry* s_i *is elliptic*

— *for all* $i, j \in \{1, \dots, n\}$, $s_i s_j$ *is elliptic.*

Corollary 2.4.6. *There is an algorithm which finds whether there exists a pseudo-periodic conjugacy class in* G.

Proof. We can algorithmically compute finitely many generating sets for the stabilizers $\mathrm{Stab}(VY(\eta))$ for finitely many pINPs η representing every pINP orbit in T.

By Lemma 2.4.1 there exist pseudo-periodic conjugacy classes if and only if one of these subgroups contain a loxodromic element.

Since a finitely generated subgroup is either elliptic or contains a loxodromic element, it then suffices to check that the algorithm returns only elliptic subgroups using Serre's lemma (Lemma 2.4.5). If it returns only elliptic subgroups then there exists no pseudo-periodic conjugacy class in G, otherwise there exists one. □

2.4.2 Link with pseudo-periodic subgroups

We now introduce *pseudo-periodic subgroups*, which can be defined independently of any train track representative. We will prove that when a train track map does exist, elements of pseudo-periodic conjugacy classes are pseudo-periodic for some ψ in the outer automorphism class of $[\phi^k]$ for some $k \in \mathbb{N}$. Pseudo-periodic subgroups which contain loxodromic elements are contained in the subgroups $\mathrm{Stab}(VY(\eta))$ defined before, and maximal pseudo-periodic subgroups containing loxodromic elements coincide with some $\mathrm{Stab}(VY(\eta))$.

Definition 2.4.7. Let $\phi \in \mathrm{Aut}(G)$ be an actual automorphism. An element $g \in G$ is *pseudo-periodic* for ϕ if for any $T \in \mathcal{D}$, for any $x \in T$, the sequence $\left(d_T(x, \phi^n(g)x)\right)_{n\in\mathbb{N}}$ is bounded.

Remarks 2.4.8.

1. The boundedness of the sequence $\left(d_T(x, \phi^n(g)x)\right)_{n\in\mathbb{N}}$ depends neither on T nor on x.
2. This notion is not defined for conjugacy classes of g. There may exist a pseudo-periodic element $g \in G$ and a conjugate of g which is not pseudo-periodic. The conjugate will be pseudo-periodic for some $\phi' \in \mathrm{Aut}(G)$ such that ϕ' and ϕ belong to the same outer class.

3. The definition also applies to elliptic elements. Not every elliptic element is pseudo-periodic in general.
4. If g is pseudo-periodic and loxodromic, then its conjugacy class is pseudo-periodic for $\Phi = [\phi]$.
5. The set $\{g \in G / g \text{ pseudo-periodic }\}$ is a subgroup of G. Indeed suppose g, h are pseudo-periodic. For every $n \in \mathbb{N}$ we have

$$\begin{aligned} d_T(x, \phi^n(gh)x) &\leq d_T(x, \phi^n(g)x) + d_T(\phi^n(g)x, \phi^n(g)\phi^n(h)x) \\ &= d_T(x, \phi^n(g)x) + d_T(x, \phi^n(h)x). \end{aligned}$$

6. A loxodromic element $g \in G$ is pseudo-periodic for ϕ if and only if there exists $m \in \mathbb{N}$ such that g^m is pseudo-periodic. The direct implication is immediate. For the converse, observe that since $\phi^n(g)$ is loxodromic for every $n \in \mathbb{N}$ we have $d_T(x, \phi^n(g)x) \leq d_T(x, \phi^n(g^m)x)$.
7. For any $n \in \mathbb{N}$ we have the inclusion $G_\phi \subset G_{\phi^n}$.

Definition 2.4.9. Let $\phi \in \mathrm{Aut}(G)$. The *pseudo-periodic subgroup* G_ϕ associated to ϕ is the subgroup of pseudo-periodic elements for ϕ.

Lemma 2.4.10. *Let $\phi \in \mathrm{Aut}(G)$ and $g \in G_\phi$ be a loxodromic pseudo-periodic element. For each $T \in \mathcal{D}$, there exists $\in\in \mathbb{N}$ such that for every $k \in \mathbb{Z}$, $\phi^k(g)$ and $\phi^{k+n}(g)$ have a pair of fundamental domains in common.*

Proof. Since $[g]$ is a pseudo-periodic conjugacy class, by Proposition 2.3.17, up to replacing ϕ by a power, we may suppose that $\|\phi^n(g)\|_T = \|g\|_T$ for every $n \in \mathbb{N}$.

Let $x \in T$. Since g is pseudo-periodic there exists $r > 0$ such that for all $n \in \mathbb{N}$, $d_T(x, \phi^n(g)x) \leq r$. Therefore, for every $n \in \mathbb{N}$, x is at distance at most $r/2$ of the axis of $\phi^n(g)$. The intersection $\mathrm{Axis}_T(\phi^n(g)) \cap \overline{B(x, 2r)}$ is a segment with endpoints in $\overline{B(x, 2r)}$ and length at least $3r$. Since there are finitely many such segments, there exist n, m with $m > n + N$ such that $\mathrm{Axis}_T(\phi^m(g)) \cap B(x, 2r) = \mathrm{Axis}_T(\phi^n(g)) \cap B(x, 2r)$. Moreover we have $\|\phi^n(g)\|_T = \|g\|_T \leq r$. These axes overlap over more than $2\|\phi^n(g)\|_T = 2\|\phi^m(g)\|_T$. This means they share at least two consecutive fundamental domains in T.

By Lemma 2.3.14, n does not depend on T: for any $S \in \mathcal{D}$, if $g, \phi^n(g)$ share two fundamental domains in T then they also do in S. □

Remark 2.4.11. For a loxodromic pseudo-periodic element $g \in G$ we can predict two behaviours for $\phi^n(g)$ depending on the *modulus* $\Delta(g) \in \mathbb{Q}^*$. The modulus is a morphism $\Delta : G \to \mathbb{Q}^*$ defined as follows; see [Lev07] for a more detailed presentation. Fix an elliptic element $a \in G$. Since the commensurator of any elliptic element of G is G, there exist $p, q \in \mathbb{Z} \setminus \{0\}$ such that $ga^pg^{-1} = a^q$. The ratio p/q does not depend on the choice of a. Define $\Delta(g) := p/q \in \mathbb{Q}$.

Suppose g is pseudo-periodic. By Lemma 2.4.10 there exists $n > 1$ such that $\phi^n(g)$ and g share a fundamental domain σ. Therefore $\phi^n(g)g^{-1}$ fixes one endpoint of σ. Let $a \in G$ be a generator of the stabilizer of this endpoint: then $\phi^n(g)g^{-1} \in \langle a \rangle$. There exists $k \in \mathbb{Z}$ such that $\phi^n(g) = a^k g$.

Let $p, q \in \mathbb{Z} \setminus \{0\}$ be such that $ga^pg^{-1} = a^q$.

— Suppose $\Delta(g) \neq 1$. Then $p - q \neq 0$ and $a^{p-q}g = a^p g a^{-p}$. Thus if $l_1 = l_2 \mod p - q$ then $a^{l_1}g$ and $a^{l_2}g$ are conjugate.

  Since there exist infinitely many possible choices for n, there exist n, n' such that $\phi^n(g) = a^k g$ and $\phi^{n'}(g) = a^{k'}g$ with $k = k' \mod p - q$. So $\phi^n(g), \phi^{n'}(g)$ are conjugates, and so are g and $\phi^{n'-n}(g)$. Thus the conjugacy class of g is actually periodic. This does not imply that there are finitely many axes among the axes of $\{\phi^n(g)/n \in \mathbb{N}\}$!

— Suppose $\Delta(g) = 1$, so $q = p$. This is the *unimodular* case. Then $a^p g = g a^p$ so a^p fixes the axis of g. There exist $\phi^n(g) = a^k g, \phi^{n'}(g) = a^{k'}g$ with $k = k' \mod p$ so $\phi^n(g)$ and $\phi^{n'}(g)$ have the same axis. By Lemma 2.3.14 this implies that g and $\phi^{n'-n}(g)$ have the same axis, and that $\phi^l(g)$ and $\phi^{l+n'-n}(g)$ have the same axis for every $l \in \mathbb{Z}$. In particular there exists $m > 1$ such that the axis of $\phi^{mn}(g)$ is the same as the axis of g for every $n \in \mathbb{Z}$. Here the axis is "periodic", but the conjugacy class of g is not in general.

In the rest of the section, we assume that $\Phi \in \mathrm{Out}(G)$ is an automorphism of G and it has a train track representative $f_0 : T \to T$. We will study automorphisms in $\mathrm{Aut}(G)$ which are in the outer automorphism class Φ^k for some $k \in \mathbb{N}$. For ψ in the outer automorphism class Φ^k, we will always use the train track representative f such that $f = g \cdot f_0^k$ for some $g \in G$, such that f is ψ-equivariant. This way, periodic Nielsen paths are the same for every $\psi \in \Phi^k$.

In our study, automorphisms whose representative has periodic points hold a special role.

Definition 2.4.12. Let $\psi \in \mathrm{Aut}(G)$ be in the outer class of Φ^k for some $k \in \mathbb{Z}$ and $f : T \to T$ be a ψ-equivariant train track map. Let η be a pINP for f. We say that ψ is *adapted to* η if the endpoints of η are f-periodic.

Remark 2.4.13. The set of pINPs for Φ is the same as for Φ^n, for any $n \in \mathbb{N}$. Thus a map f is adapted to η if and only if there exists $n \in \mathbb{N}$ such that f^n is adapted to η.

Proposition 2.4.14. *Let k be such that for every pINP η in T, $[f_0^k(\eta)]$ is in the orbit of η. For every loxodromic $g \in G$ such that the conjugacy class $[g]$ is pseudo-periodic for Φ, there exists $\psi \in \Phi^k$ with an associated train track representative f such that $g \in G_\psi$ and (ψ, f) is adapted to a pINP η in the decomposition of the axis of g given by Proposition 2.3.18.*

Therefore for every loxodromic g in a pseudo-periodic conjugacy class $[g]$ of G, there exists a pINP η and a pair (ψ, f) adapted to η such that $g \in G_\psi$. Moreover $g \in \mathrm{Stab}(VY(\eta))$.

Proof. Let $[g]$ be a pseudo-periodic conjugacy class. By Proposition 2.3.18 the axis of g is a concatenation of pINPs. Take a pINP $\eta = [x, y]$ in $\mathrm{Axis}_T(g)$: there exists $h \in G$ such that $[f_0^k(\eta)] = h\eta$. Define $\psi := c_{h^{-1}} \circ \phi^k$. The map $f := h^{-1} \cdot f_0^k$ is a train track representative for ψ and the endpoints x, y of η are fixed points for f.

For every $n \in \mathbb{N}$, $f^n(x) = x$. Since $[x, gx]$ is a concatenation of pINPs, the distance $d_T(f^n(x), f^n(gx)) = d_T(x, \phi^n(g)x)$ is bounded so $g \in G_\psi$. □

A point $x \in T$ is a *pre-periodic point* for f if there exists $n \geq 0$ such that $f^n(x)$ is periodic. We say x is *non-escaping* if for any $y \in T$ the distance $d_T(y, f^n(x))$ is bounded for $n \to \infty$. A pre-periodic point is obviously non-escaping but the converse is also true for a train track map:

Lemma 2.4.15. *Let $\phi \in \mathrm{Aut}(G)$ and $f : T \to T$ be a train track representative for ϕ. If x is non-escaping then there exists $n \in \mathbb{N}$ and $m \in \mathbb{N} \setminus \{0\}$ such that $f^{n+m}(x) = f^n(x)$, so x is actually a pre-periodic point.*

Proof. Suppose for some $n_0 \in \mathbb{N}$, $f^{n_0}(x)$ is a vertex. Since f maps vertices to vertices and T is locally finite, $f^n(x)$ takes only finitely many values so there exists $n < m$ such that $f^n(x) = f^m(x)$ and $f^n(x)$ is an actual periodic point.

If $f^n(x)$ is in the interior of an edge for every $n \in \mathbb{N}$, there exists $n < m$ such that $f^n(x)$ and $f^m(x)$ are both contained in the same edge e. The image $f^{m-n}(e)$ contains e so there exists a fixed point $x_0 \in e$ for f^{m-n}. Since e is a legal segment, it is stretched

uniformly by a factor $\lambda > 1$ and unless $f^n(x) = x_0$, the distance $d_T(x_0, f^{k(m-n)}(x))$ grows exponentially when $k \to +\infty$. Therefore $f^n(x) = x_0$. □

Suppose that the pair (ϕ, f) is adapted to a pINP $\eta \subset T$. The points of $VY(\eta)$ are non-escaping so they are also pre-periodic for f.

The following lemma is a general result about actions on trees.

Lemma 2.4.16. *Let H be a group acting on a simplicial tree T by isometries such that there exist loxodromic isometries. Suppose that there exists an elliptic element $a \in H$ such that for every loxodromic $h \in H$, the product ah is an elliptic isometry. Then the action of H on T is dihedral.*

Proof. As in [GL07] we distinguish 4 possible types of action for H: linear abelian, dihedral, genuine abelian and irreducible. We prove the lemma by contraposition, i.e. if the action is abelian or irreducible then for every elliptic a there exists a loxodromic $h \in H$ such that ah is loxodromic.

Suppose that the action is abelian (linear or genuine). It has a fixed point $\xi \in \partial T$. Let $a \in H$ be an elliptic isometry. Let $h \in H$ be a loxodromic isometry. Both h and a fix ξ so there exists $x \in \mathrm{Axis}_T(h)$ such that the subray $[x, \xi]$ is contained in $\mathrm{Fix}_T(a)$. The isometry ah acts on $[x, \xi]$ like h so it is also loxodromic.

Otherwise suppose that the action is irreducible. Let $a \in H$ be an elliptic element, and $h \in H$ be a loxodromic element. Suppose $\mathrm{Axis}_T(h) \cap \mathrm{Fix}_T(a) = \varnothing$. Let us prove that ah is loxodromic. Let $p \in T$ be the projection of $\mathrm{Fix}_T(a)$ on $\mathrm{Axis}_T(h)$. The point $p' := h^{-1}p$ is sent by ah to the point ap. Let $v \in T$ be a point in $[p', p]$. Then $ahv \in [ap, ahp]$. In particular ahv is not in $[p', ahp']$ so ah is loxodromic and p' belongs to its axis.

Suppose now that $\mathrm{Axis}_T(h) \cap \mathrm{Fix}_T(a)$ is non empty. If a fixes one of the endpoints of $\mathrm{Axis}_T(h)$ then the same discussion as for the abelian case applies. Else there exists a loxodromic element $g \in H$ whose axis does not intersect the axis of h. Since $\mathrm{Axis}_T(h)$ is not fixed by a, there exists $k \in \mathbb{Z}$ such that $h^k \mathrm{Axis}_T(g)$ does not intersect $\mathrm{Fix}_T(a)$ so up to replacing g by the conjugate $h^k g h^{-k}$ we may assume $\mathrm{Axis}_T(g) \cap \mathrm{Fix}_T(a) = \varnothing$. Then ag is loxodromic. □

Suppose that a pair (ϕ, f) is adapted to *eta*. The following is the key for proving that $VY(\eta)$ is stable by G_ϕ:

Lemma 2.4.17. *Suppose the group G_ϕ of pseudo-periodic elements contains a loxodromic element. Let $\phi \in \mathrm{Aut}(G)$ with a train track representative $f : T \to T$. Suppose that (ϕ, f) is adapted to $\eta = [x, x']$ a pINP in T. Then $G_\phi \subset \mathrm{Stab}(VY(\eta))$.*

Proof. Suppose first that g is a loxodromic element. By Remark 2.4.13, f is adapted to η if and only if its powers are. By Remark 2.4.8 $G_\phi \subset G_{\phi^n}$ so it suffices to prove the lemma for the pair (ϕ^n, f^n) for some $n \in \mathbb{N}$. Thus we may assume for simplicity that $f(x) = x$.

By Lemma 2.4.1 the points gx and g^2x are pre-periodic. They might not be periodic points. If gx is a periodic point, the path $[x, gx]$ is a periodic Nielsen path so by Lemma 2.3.2 it is a concatenation of pINP.

Suppose otherwise. The axis of g is a legal concatenation of pINPs, by Proposition 2.3.18. We will prove that all these pINPs belong to $EY(\eta)$. Then by Lemma 2.4.1 $g \in \mathrm{Stab}(VY(\eta))$. This will conclude the proof for g loxodromic.

See Figure 2.6 for a picture of this proof.

The point x is fixed by f and since the path between x and g^2x is a concatenation of pINPs, the point g^2x is non-escaping, so it is pre-periodic. Thus there exists $N \in \mathbb{N}$ such that for all $n \geq N$, $f^n(g^2x) = \phi^n(g^2)x$ is a periodic point. By Lemma 2.4.10 we may choose n such that $\phi^n(g^2)x$ is periodic and the axes of $\phi^n(g)$ and g intersect along two common fundamental domains. By Lemma 2.4.1 we have $\phi^n(g^2) \in \mathrm{Stab}(VY(\eta))$.

The path $[x, \phi^n(g^2)x]$ is a concatenation of pINPs. It contains a pINP $\eta' \subset \mathrm{Axis}_T(\phi^n(g))$. By Lemma 2.3.2 η' appears in the decomposition of the Nielsen path $[x, \phi^n(g^2)x]$. Thus $\eta' \in EY(\eta)$.

Since $\phi^n(g)$ and g share two fundamental domains, there exists $l \in \mathbb{Z}$ such that $\phi^n(g)^l\eta' \subset \mathrm{Axis}_T(g) \cap \mathrm{Axis}_T(\phi^n(g))$. By Lemma 2.3.2 this translate of η' appears in the pINP decomposition of $\mathrm{Axis}_T(g)$, besides it belongs to $EY(\eta)$. This completes the proof when g is loxodromic.

If g is elliptic the argument using the axes does not work any more. If there exist loxodromic elements h_1, h_2 such that $g = h_1h_2$ then the path $[x, gx]$ can be written as the concatenation $[x, h_1x] \cdot h_1[x, h_2x]$ and the loxodromic case completes the proof.

If g cannot be written as a product of loxodromic elements then by Lemma 2.4.16 the action of G_ϕ on its minimal subtree is dihedral (by assumption it is not elliptic). Let ℓ be the invariant axis in T. If g fixed ℓ then for any loxodromic element $h \in G_\phi$, gh would be loxodromic so $g = (gh)h^{-1}$ would be a product of loxodromic elements. Thus g acts like a symmetry. There exists a loxodromic isometry $h \in G_\phi$ with axis ℓ. By Proposition 2.3.18 ℓ is a legal concatenation of pINPs.

As in the loxodromic case, the pINPs in the decomposition of $\ell = \mathrm{Axis}_T(h)$ belong to $EY(\eta)$. Furthermore this decomposition into pINPs is g-invariant. Let $y \in VY(\eta)$ be an endpoint of a pINP in ℓ, then the point gy is also in $EY(f)$. By Lemma 2.4.1 we have

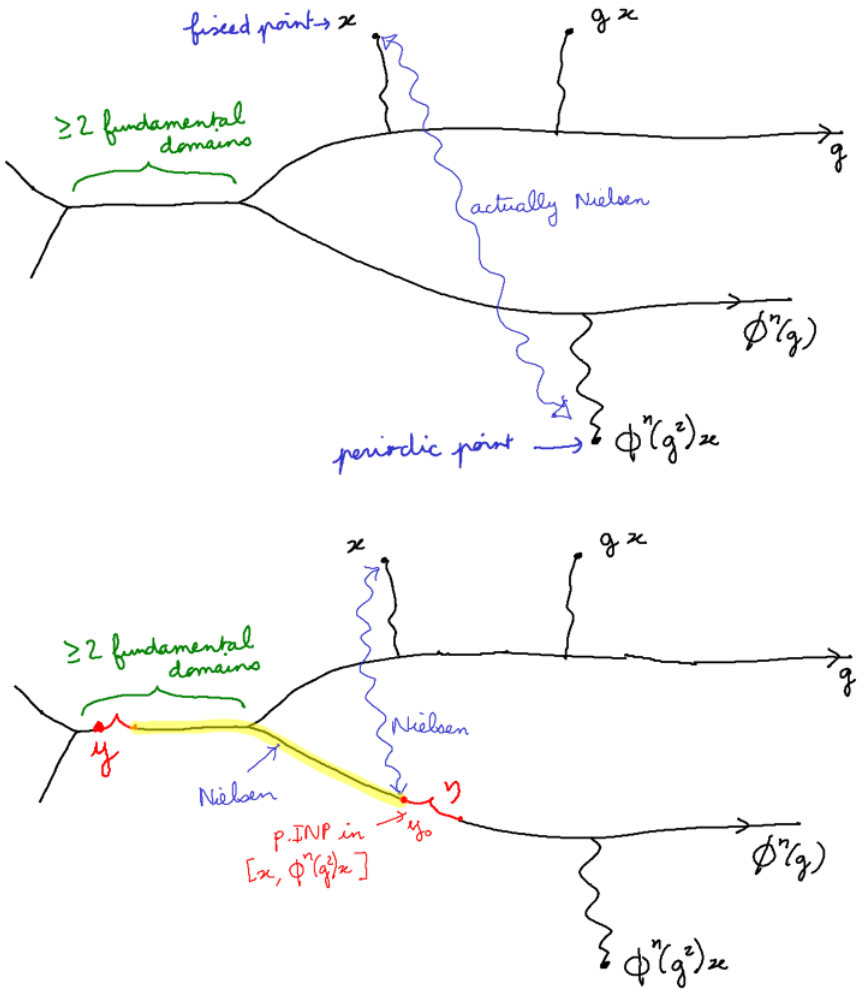

Figure 2.6 – Finding a concatenation of pINPs between x and gx

$g \in \mathrm{Stab}(VY(\eta))$. This concludes the proof in the elliptic case. □

Corollary 2.4.18. *Let $\phi \in \mathrm{Aut}(G)$ and let f be a train track representative for ϕ. Suppose that (ϕ, f) is adapted to η. Then $\mathrm{Stab}(VY(\eta)) = G_\phi$.*

Proof. The inclusion $G_\phi \subset \mathrm{Stab}(VY(\eta))$ is proved by Lemma 2.4.17. Let us prove the inverse inclusion. Suppose $g \in \mathrm{Stab}(VY(\eta))$. Let x_0 be an endpoint of η. Then $gx_0 \in VY(\eta)$ so gx_0 is pre-periodic. For all $n \in \mathbb{N}$ we have $f^n(gx_0) = \phi^n(g)f^n(x_0)$. Since x_0 and gx_0 are pre-periodic, there exists $B > 0$ such that the sets $\{f^n(x_0), n \in \mathbb{N}\}$ and $\{\phi^n(g)f^n(x_0), n \in \mathbb{N}\}$ have diameter smaller than B. Then for all $n \in \mathbb{N}$ we have

$$d_T(x_0, \phi^n(g)x_0) \leq d_T(x_0, \phi^n(g)f^n(x_0)) + \phi^n(g)d_T(f^n(x_0), x_0) \leq 2B$$

so $g \in G_\phi$. □

Pseudo-periodic subgroups defined by pairs (ϕ, f) which are adapted to a pINP have good properties but not all pseudo-periodic subgroups arise in this way. However we will see that it is the case for *maximal* pseudo-periodic subgroups.

Proposition 2.4.19. *Let $\psi \in [\phi^k]$ for some $k \in \mathbb{N}$, and let f be the associated train track representative. If the action of G_ψ on T is irreducible then the pair (ψ, f) is adapted to a pINP.*

Proof. Suppose that the action of G_ψ is irreducible. Then we claim that G_ψ contains unimodular loxodromic elements $u, v \in G$ whose axes have distinct endpoints and an intersection longer than $4C_f$. We postpone the proof of the claim.

By Remark 2.4.11 there exists $n \geq 1$ such that u and $\phi^n(u)$ have the same axis with same orientation, and have the same action on it. Recall that f^n sends maximal legal segments to legal segments. Writing the axis of u as a bi-infinite concatenation of pINPs $\{\eta_i, i \in \mathbb{Z}\}$, f^n induces a translation on the set of pINPs: there exists $k_u \in \mathbb{Z}$ such that $[f^n(\eta_i)] = \eta_{i+k_u}$.

Similarly f^n shifts the pINPs $\{\eta'_j, j \in \mathbb{Z}\}$ in $\mathrm{Axis}_T(v)$ by $k_v \in \mathbb{Z}$. The length of a pINP is at most $2C_f$ so the intersection $\mathrm{Axis}_T(v) \cap \mathrm{Axis}_T(u)$ contains a pINP $\eta_i = \eta'_j$ belonging to the decomposition of both axes. Suppose by contradiction that $k_u \neq 0$ is non-zero, then there exists $m \in \mathbb{N}$ such that $[f^{nm}(\eta_i)] \cap \mathrm{Axis}_T(u) = \eta_{i+mk_u}$, which is non-empty, does not intersect $\mathrm{Axis}_T(v)$. However $[f^{nm}(\eta_i)] = f^{nm}(\eta'_j) = \eta'_{j+mk_v} \subset \mathrm{Axis}_T(v)$. This is a contradiction.

Thus $f^n(\eta_i) = \eta_i$ so (ψ, f) is adapted to η_i.

Now let us prove the claim. We will use Lemma 2.4.20 below to construct u and v as announced. Since the action of G_ψ is irreducible there exist loxodromic elements $g, h \in G_\psi$ with disjoint axes. The axes of g and $hg^{-1}h^{-1}$ are disjoint so the element $u := ghg^{-1}h^{-1}$ is unimodular and loxodromic, and $\|u\|_T \geq 2\|g\|_T$. Up to replacing g by a power we may suppose $\|u\|_T \geq 4C_f$. Moreover the axes of g and u intersect along one fundamental domain of g, so there exists $l \in \mathbb{Z}$ such that the axes of u an $g^l u^{-1} g^{-l}$ do not intersect. Thus the element $v := g^l u^{-1} g^{-l} u$ is also unimodular and loxodromic, moreover the axes of u and v overlap on a length equal to $\|u\|_T$, which proves the claim. □

Lemma 2.4.20. *Let T be a simplicial tree. Let a, b be loxodromic isometries of T with disjoint axes. Then the product ab is a loxodromic isometry with translation length $\|a\|_T +$*

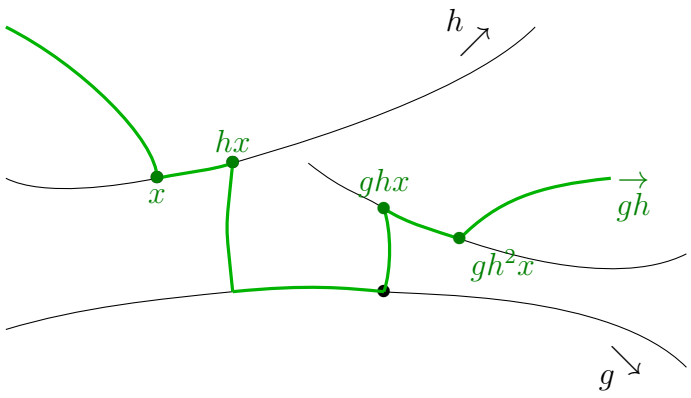

Figure 2.7 – Axis of gh when g, h have disjoint axes

$\|b\|_T + 2d_T(\mathrm{Axis}_T(a), \mathrm{Axis}_T(b))$. *Moreover* $\mathrm{Axis}_T(ab) \cap \mathrm{Axis}_T(a)$ *is a single fundamental domain for* a.

Proof. Most of this is proved in [CM87, p. 1.5]; we give the proof for completeness but it is best explained by Figure 2.7. Let a, b be as above. Let p_a be the projection of $\mathrm{Axis}_T(b)$ to $\mathrm{Axis}_T(a)$, and p_b be the projection of $\mathrm{Axis}_T(a)$ to $\mathrm{Axis}_T(b)$. Let $x := b^{-1}p_b$. Then $abx = ap_b$ and the path $[x, abx]$ is the union $[x, p_b] \cup [p_b, p_a] \cup [p_a, ap_a] \cup [ap_a, ap_b]$. The overlap of those 4 segments has zero length because p_a, p_b are projections. The path $[p_b, abp_b] = [p_b, p_a] \cup [p_a, ap_a] \cup [ap_a, ap_b] \cup [ap_b, abp_b]$ has the same length since $[ap_b, abp_b] = ab[x, p_b]$ and the decomposition also has trivial overlap. If x were not in the axis, the point p_b which lies in the interior of $[x, abx]$ would be moved by a shorter distance, so x is actually in the axis. The translation length is $\|a\|_T + \|b\|_T + 2d_T(p_a, p_b)$ as announced and $\mathrm{Axis}_T(ab) \cap \mathrm{Axis}_T(a) = [p_a, ap_a]$. □

Pseudo-periodic subgroups containing loxodromic elements are contained in stabilizers of Nielsen classes, with equality when the subgroup comes from a pair adapted to a pINP:

Lemma 2.4.21. *For every* ψ *in* Φ^k *for* $k \in \mathbb{N}$*, if* G_ψ *contains a loxodromic element, there exists* $\eta \subset T$ *a pINP such that* $G_\psi \subset \mathrm{Stab}(VY(\eta))$.

If (ψ, f) *is adapted to a pINP then the inclusion is an equality.*

Proof. The second statement, for (ψ, f) adapted to a pINP η, is a consequence of Corollary 2.4.18: we have $G_\psi = \mathrm{Stab}(VY(\eta))$.

Suppose G_ψ contains a loxodromic element. If the action of G_ψ on T is irreducible then by Proposition 2.4.19 (ψ, f) is adapted to a pINP η so the second statement applies.

Otherwise suppose that there is a fixed point in ∂T for the action of G_ψ. Let g be such that the attracting point for g is fixed by G_ψ.

By Proposition 2.4.14 there exists a pINP $\eta \subset \mathrm{Axis}_T(g)$ and a pair (ψ', f') adapted to η so $g \in \mathrm{Stab}(VY(\eta)) = G_{\psi'}$. Let $h \in G_\psi$ be a loxodromic element. The axes of g and h intersect along an infinite ray. There exist powers l, m such that $\|g^l\|_T = \|h^m\|_T$ so g^l and h^m have at least two fundamental domains in common, and by Lemma 2.3.14 for every ϕ in the outer class of Φ and every $n \in \mathbb{Z}$ the elements $\phi^n(g^l)$ and $\phi^n(h^m)$ also share two fundamental domains.

Up to replacing η by some translate $g^j\eta$, we may assume $\eta \subset \mathrm{Axis}_T(h)$ so $g^l h^{-m}\eta = \eta$. Since the conjugacy class $[h]$ is pseudo-periodic, $\mathrm{Axis}_T(h)$ is a concatenation of pINPs and by Lemma 2.3.2, η is a pINP of the decomposition. Thus the pINPs in $\mathrm{Axis}_T(h)$ belong to $EY(\eta)$. By Lemma 2.4.1 $h \in \mathrm{Stab}(VY(\eta))$.

By Lemma 2.4.16 all elliptic elements can be written as the product of loxodromic elements so since all loxodromic elements of G_ψ belong to $\mathrm{Stab}(VY(\eta))$, we have $G_\psi \subset \mathrm{Stab}(VY(\eta))$.

The only remaining case is when the action of G_ψ is dihedral. In that case G_ψ preserves an axis ℓ in T.

There exists a loxodromic element $g \in G_\psi$ whose axis is ℓ. Let η be a pINP in ℓ: then the endpoints of the pINPs in the decomposition of ℓ belong to $VY(\eta)$. Since this decomposition is unique by Lemma 2.3.2 and G preserves the set of pINPs in T, the group G_ψ preserves $VY(\eta) \cap \ell$. By Lemma 2.4.1 this implies $G_\phi \subset \mathrm{Stab}(VY(\eta))$. □

Thus maximal pseudo-periodic subgroups which contain loxodromic elements come from pairs adapted to a pINP:

Proposition 2.4.22. *A pseudo-periodic subgroup H containing a loxodromic element is maximal for inclusion among pseudo-periodic subgroups if and only if there exists a pINP $\eta \subset T$ and a pair (ψ, f) adapted to η such that $H = G_\psi$.*

Proof. Suppose H is maximal. By Lemma 2.4.21 there exists a pINP η such that $H \subset \mathrm{Stab}(VY(\eta))$. There exists $\psi \in \Phi^k$ for some $k \in \mathbb{N}$ and $g \in G$ such that $f := gf_0^k$ fixes the endpoints of η, so the pair (ψ, f) is adapted to η and $\mathrm{Stab}(VY(\eta)) = G_\psi$. By maximality $H = G_\psi$.

Conversely suppose there exists a pair (ψ, f) such that $H = G_\psi$ adapted to a pINP $\eta \subset T$. By Corollary 2.4.18 $H = \mathrm{Stab}(VY(\eta))$. Let H' be a pseudo-periodic subgroup

such that $H \subset H'$: by Lemma 2.4.21 there exists η' such that $H' \subset \mathrm{Stab}(VY(\eta'))$. Thus for $h \in H$, $\mathrm{Axis}_T(h) \subset \mathrm{conv}(VY(\eta'))$. By uniqueness of the decomposition of $\mathrm{Axis}_T(h)$ into pINPs, the pINPs in h belong to $EY(\eta)$ and $EY(\eta')$ so $VY(\eta) = VY(\eta')$ and $H' \subset \mathrm{Stab}(VY(\eta)) = H$, hence the maximality of H. □

Corollary 2.4.23. *There are finitely many conjugacy classes of maximal pseudo-periodic subgroups containing loxodromic elements associated to the outer classes $\Phi^k \in \mathrm{Out}(G)$ for $k \in \mathbb{N}$. Furthermore, these subgroups are finitely generated.*

Proof. This is a consequence of Proposition 2.4.22 and Corollary 2.4.18: the maximal pseudo-periodic subgroups containing loxodromic elements coincide with the stabilizers of Nielsen classes containing loxodromic elements.

The conclusion follows from the fact that there are finitely many orbits of Nielsen classes and their stabilizers are finitely generated. □

2.5 Algorithm

In this section we assemble the results obtained before to give the desired algorithm. First let us explain how to deal with restricted deformation spaces.

Let G be a GBS group represented by a graph of groups Γ. As in Chapter 1 we say that a family of subgroups $\mathcal{A}$ which is invariant by conjugacy and taking subgroups is *represented* by finite sets of integers $(I_v)_{v \in V(\Gamma)}$ if for every $v \in \Gamma$ and every lift $\tilde{v}$ in the universal cover of Γ, the set I_v is the set of minimal elements of $\{[G_v : G_v \cap A], A \in \mathcal{A}\}$ for divisibility.

Theorem 2.5.1. *There is an algorithm which takes*

— *a non-elementary GBS group G*

— *an automorphism $\phi \in \mathrm{Aut}(G)$*

— *a train track representative $f : T \to T$ for ϕ where T is given as a graph of groups*

which finds out whether ϕ is pseudo-atoroidal.

Proof. This is a consequence of Corollary 2.4.6. □

Theorem 2.5.2. *There is an algorithm which takes*

— *a non-elementary GBS group G*

— *a pseudo-atoroidal automorphism* $\phi \in \mathrm{Aut}(G)$

— *a train track representative* $f : T \to T$ *for* ϕ *where* T *is given as a graph of groups*

— *a family of sets* $(I_v)_{v \in V(T/G)}$ *representing the family of subgroups* $\mathcal{A}$

which finds out whether ϕ *has no pseudo-periodic element. In that case it decides whether* ϕ *is iwip for* $\mathcal{D}^{\mathcal{A}}$.

In the previous theorem, ϕ need not be pseudo-atoroidal as long as it has no simple pseudo-periodic conjugacy class. Note that we do not know how to check this condition algorithmically.

Proof. There are two steps in this algorithm.

The first step consists in finding either a primitive train track representative or a proof of reducibility for ϕ. Corollary 2.1.12 solves this problem algorithmically using the matrix $A(f)$ and produces a primitive train track representative if it exists.

The second step consists in computing turns of the stable lamination and its Whitehead graphs, then using Corollary 2.2.7. Using Lemma 2.1.21 we can compute all Whitehead graphs. Then compute connected components of $\mathrm{Wh}_T(\Lambda^+, v)$ for every $v \in V(T/G)$. Let W be such a Whitehead graph. If W is not connected we need to compute the stabilizers of connected components. As in Section 1.3 of Chapter 1, one can compute for every connected component C the index $i(C) := [G_v : \mathrm{Stab}(C)]$. Then $\mathrm{Stab}(C) \in \mathcal{A}$ if and only if $i(C)$ is divisible by some $i \in I_v$.

If there exists a Whitehead graph with a connected component whose stabilizer is in $\mathcal{A}$ then by Proposition 2.2.3 ϕ is reducible. Moreover the proof of the proposition also gives an invariant class of special factors with respect to $\mathcal{A}$ for ϕ.

Otherwise Corollary 2.2.7 states that ϕ is fully irreducible. □

Chapitre 3

Strongly contracting axes for fully irreducible automorphisms of Generalized Baumslag-Solitar groups

Introduction

Let G be a generalized Baumslag-Solitar (GBS) group, i.e. a group which is isomorphic to the fundamental group of a graph of groups where all vertex and edge groups are infinite cyclic. By Bass-Serre theory, G admits a minimal action by isometries on a simplicial tree T such that all vertex or edge stabilizers in T are infinite cyclic.

In general G admits infinitely many such actions. When G is not isomorphic to $\mathbb{Z}, \mathbb{Z}^2$ or the fundamental group of a Klein bottle $\langle a, b | a^2 = b^2 \rangle = \langle a, t | tat^{-1} = a^{-1} \rangle$, then the *cyclic deformation space* $\mathcal{D}$ associated to G is defined as the projectivized set of minimal actions of G by isometries on metric simplicial trees with edge and vertex stabilizers isomorphic to $\mathbb{Z}$, where actions T, T' are identified if there exists a G-equivariant isometry $T \to T'$. The space $\mathcal{D}$ is analogous to Culler and Vogtmann's Outer Space CV_N for the free group F_N. The study of Outer Space is crucial for understanding the outer automorphism group $\mathrm{Out}(F_N)$. Just as $\mathrm{Out}(F_N)$ acts on CV_N, the outer automorphism group $\mathrm{Out}(G)$ acts on $\mathcal{D}$ by pre-composition of the action: if G acts on T, then define $T \cdot \phi$ as the action whose underlying space is T and where for $t \in T$, $g \cdot_{T \cdot \phi} t = \phi(g) \cdot_T t$.

There is an important analogy between the study of the outer automorphism group of free groups F_N and mapping class groups. In this analogy, Outer Space is the counterpart for Teichmüller space. Fully irreducible automorphisms of $\mathrm{Out}(F_N)$ are a special class of automorphisms which do not preserve the conjugacy class of any free factor. In some sense they act on Outer Space by translations along an axis whose points are actions

which admit a *train track representative*. Their equivalent in the mapping class groups context is pseudo-Anosov automorphisms. Such automorphisms can also be defined for GBS groups: we say that an automorphism $\phi \in \mathrm{Out}(G)$ is *fully irreducible* if no conjugacy class of proper *cyclic factor* is ϕ-periodic. A *cyclic factor* is the analogue of a free factor for the free group: it is a non-cyclic subgroup $G' < G$ such that there exists a graph of groups Γ representing G, and a subgraph $\Gamma' \subset \Gamma$ such that the fundamental group $\pi_1(\Gamma')$ identifies to the subgroup G'. We say that G' is a proper cyclic factor if $G' \neq G$.

One main property for a fully irreducible automorphism $\phi \in \mathrm{Out}(F_N)$ is that its action on CV_N is hyperbolic and admits an axis $\mathcal{L}_\phi$. In order for this to make sense, we endow CV_N with the non-symmetric *Lipschitz metric* d_{Lip}: it is defined by

$$d_{\mathrm{Lip}}(T, T') := \mathrm{Lip}(T, T') \frac{\mathrm{vol}(T)}{\mathrm{vol}(T')}$$

for $T, T' \in \mathrm{CV}_N$, where $\mathrm{Lip}(T, T') := \sup_{f:T \to T' \text{ Lipschitz}} \mathrm{Lip}(f)$.

One can define a closest point projection $\pi_f : \mathcal{D} \to \mathcal{L}_f$ such that for $T \in \mathcal{D}$, $d_{\mathrm{Lip}}(T, \pi_f(T))$ is minimal. Define the *outward ball* $B_\to(Y, r) := \{T \in \mathcal{D} / d_{\mathrm{Lip}}(Y, T) < r\}$.

In [Alg11] Algom-Kfir proves that the projections on axes have the *strong contraction property*: there exists $B > 0$ depending only on N and ϕ such that the diameter of the projection of any outward ball in $\mathcal{D}$ disjoint from $\mathcal{L}_\phi$ is bounded by B.

Likewise we can define these objects for the deformation space $\mathcal{D}$ of a GBS group. A few differences arise, for example because trees in $\mathcal{D}$ have non-trivial edge stabilizers: one consequence is that in general there exist $T, T' \in \mathcal{D}$ such that $d_{\mathrm{Lip}}(T, T') = 0$ and $d_{\mathrm{Lip}}(T', T)$ is arbitrarily large.

In this paper we prove the analogue of this property for fully irreducible automorphisms of a GBS group:

Theorem A. *Let G be a GBS group with $b_1(G) \geq 3$. Let ϕ be an atoroidal fully irreducible automorphism such that ϕ, ϕ^{-1} both admit train track representatives. Let $\mathcal{L}_\phi$ be an axis for ϕ in $\mathcal{D}$ and let π_ϕ be a closest point projection to $\mathcal{L}_\phi$. Then there exists $D > 0$ such that for any $Y \in \mathcal{D}$ and $r > 0$ such that $B_\to(Y, r) \cap \mathcal{L}_\phi = \varnothing$*

$$\mathrm{diam}(\pi_\phi(B_\to(Y, r))) \leq D$$

The proof for this result is similar to [Alg11].

We do not know whether train track representatives always exist for fully irreducible

automorphism. Several cases are known:

— The procedure in [BH92] may be adapted to construct a train track map by hand on a specific example. See Example 3.1.11.

— If G has no non-trivial integer modulus then [For06] proves that $\mathcal{D}$ has finite dimension. Then [Mei15] implies that all fully irreducible automorphisms admit train track representatives.

— If $G := \mathrm{BS}(p, pn)$ then Bouette proved in [Bou16] that all automorphisms of G are reducible and preserve the conjugacy class of a common cyclic factor H. She then introduces a new deformation space $\mathcal{D}_H$ consisting in all G-trees with cyclic edge stabilizers, and vertex stabilizers either cyclic or conjugate to H. In this deformation space, there exist fully irreducible automorphisms and they admit train track representatives.

The second case is not relevant here since we need the trees in the deformation space to be locally finite.

For technical reasons, we ask that the automorphism ϕ in the theorem be pseudo-atoroidal, which means that for all $g \in G$, $(\|\phi^n(g)\|_T)_{n\in\mathbb{N}}$ is unbounded.

We also need the first Betti number $b_1(\Gamma)$ for any graph of cyclic groups Γ with $\pi_1(\Gamma) \simeq G$ to be at least 3. Actually $b_1(\Gamma)$ does not depend on the choice of Γ when G is not isomorphic to the fundamental group of a Klein bottle, which we exclude.

An important result in [Alg11] is the fact that axes of automorphisms have bounded projection on each other. This fact is of great interest since it enables the definition of a projection complex on which the quasi-tree construction of [BBF15] could be applied. However we do not know yet if this still holds in the GBS context, due to the fact that the Lipschitz metric is even less symmetric. For example it is not true that the Lipschitz metric is quasi-symmetric on the θ-thick part of $\mathcal{D}$. Furthermore, there is no bound on the number of candidates in trees of $\mathcal{D}$.

In Section 3.1 we give some background about GBS groups and their automorphisms. We develop the topic of laminations in Section 3.2. In Section 3.3 we prove the analogue of results from [Alg11] which state that the axis of a simple element cannot follow both the stable and unstable lamination for a long distance; the method of the proof differs somehow from the original. Section 3.4 develops the behaviour of lines in T such as axes of elements of g when iterating a train track $f : T \to T$, and Section 3.5 relies on it

to define a projection of $\mathcal{D}$ on the axis of a fully irreducible element with a train track representative. The contents in this section are really close to [Alg11] and are there for completeness.

We prove negative curvature properties of the projection in Section 3.6, that is, inequalities about distances in $\mathcal{D}$. Although the former does not differ from the free group case, the latter needs some arguments which are specific to GBS groups. Finally we prove the strong contraction for balls of outward radius.

3.1 Generalities

3.1.1 Graphs and trees

A *graph* Γ is defined by $(V(\Gamma), E(\Gamma), \bar{\cdot}, o, t)$ where

— $V(\Gamma)$ is a set of *vertices*

— $E(\Gamma)$ is a set of *edges*

— the map $\bar{\cdot}$ is an involution $E(\Gamma) \to E(\Gamma)$ without fixed point; for $e \in E(\Gamma)$ the edge $\bar{e}$ is called the *opposite edge*

— the maps $o, t : E(\Gamma) \to V(\Gamma)$ are the *initial vertex* and *terminal vertex* maps, with the property that every $e \in E(\Gamma)$ satisfies $o(e) = t(\bar{e})$.

It is finite if $V(\Gamma), E(\Gamma)$ are finite. See [Ser77] for more details on graphs.

An edge path in Γ is a sequence $e_1, \dots, e_k$ with $e_i \in E(\Gamma)$ for $i \in \{1, \dots, k\}$ and $t(e_i) = o(e_{i+1})$ for $i \leq k-1$. It is *non-backtracking* if for all $i \in \{1, \dots, k-1\}$, $\bar{e}_i \neq e_{i+1}$. It is a *loop* if $o(e_1) = t(e_k)$.

A *tree* is a graph without non-backtracking loops.

Let $E^+(\Gamma)$ be an orientation of the edges, i.e. a subset of $E(\Gamma)$ such that $E^+(\Gamma) \sqcup \overline{E^+(\Gamma)}$ is a partition of $E(\Gamma)$.

A *metric* on a graph Γ is a map $\mathrm{len}_\Gamma : E(\Gamma) \to \mathbb{R}_+$ such that for all $e \in E(\Gamma)$, $\mathrm{len}_\Gamma(\bar{e}) = \mathrm{len}_\Gamma(e)$.

The *geometric realization* of a graph Γ is the union of points $(x_v)_{v \in V(\Gamma)}$ and segments $(\sigma_e)_{e \in E^+(\Gamma)}$ where σ_e is isometric to $[0, \mathrm{len}_\Gamma(e)]$ for every $e \in E(\Gamma)$, where for every $e \in E^+(\Gamma)$ we identify the first point of σ_e with $x_{o(e)}$ and its last point with $x_{t(e)}$. It is endowed with the associated path metric. It does not depend on the choice of $E^+(\Gamma)$.

In the rest of the paper we will identify trees and other graphs with their geometric realizations. A *path* in a tree T is the image of a Lipschitz map from an interval to T. It

is *non-backtracking* if the map is an immersion, and equivalently if it is a the image of a geodesic. In the context of geometric realizations an *edge path* is a path which is the image of an edge path in the graph. For two points $x, y \in T$, the segment $[x, y]$ is the unique geodesic from x to y.

A *graph of groups* is a graph Γ together with collections of vertex groups $(G_v)_{v\in V(\Gamma)}$ and edge groups $(G_e)_{e\in E(\Gamma)}$ and monomorphisms $\iota_e : G_e \to G_{t(e)}$. Let τ be a maximal subtree in the graph Γ. The fundamental group $\pi_1(\Gamma, \tau)$ of the graph of groups Γ is defined as follows:

$$\pi_1(\Gamma, \tau) = \left\langle \bigcup_{v\in V(\Gamma)} G_v \cup (t_e)_{e\in E(\Gamma)} | \bigcup_{v\in V(\Gamma)} R_v, \bigcup_{e\in E(\Gamma)} R_e, R_\tau \right\rangle$$

where

— for $v \in V(\Gamma)$, R_v is the set of relations of G_v

— for $e \in E(\Gamma)$, $R_e = \{t_e \phi_e(h) t_{\bar{e}} \phi_{\bar{e}}(h)^{-1} / h \in G_e\}$

— $R_\tau := \{t_e, e \in \tau\}$

Note that for $h = 1$ we obtain the relation $t_{\bar{e}} = t_e^{-1}$. Up to isomorphism the fundamental group does not depend on the choice of τ. If the vertex groups and edge groups are finitely generated then $\pi_1(\Gamma)$ is finitely presented.

Let G be a group. A *marked graph of groups* for G is a graph of groups Γ together with a *marking* (i.e. an identification) $\Psi : G \to \pi_1(\Gamma)$. The automorphism group of G acts as follows on the set of marked graphs of groups: if $\phi \in \mathrm{Aut}(G)$ and (Γ, Ψ) is a marked graph of groups then $(\Gamma, \Psi) \cdot \phi := (\Gamma, \Psi \circ \phi)$.

A *G-tree* T is a metric simplicial tree with an action of G by isometries. The tree T is *minimal* if there is no proper G-invariant subtree.

The *universal cover* of a graph of groups Γ is a minimal $\pi_1(\Gamma)$-tree T such that T/G is isomorphic to Γ as a graph and for every $v \in V(\Gamma)$ and every lift $\bar{v} \in T$, the stabilizer of $\bar{v}$ is isomorphic to G_v. By Bass-Serre theory in [Ser77], universal covers exist and are unique up to $\pi_1(\Gamma)$-equivariant isomorphism.

Moreover Bass-Serre theory gives a correspondance between marked graphs of groups for G and G-trees.

If Γ is a metric graph then the metric naturally lifts to its universal cover.

In a G-tree T we denote the pointwise stabilizer of a vertex v (resp. an edge e) by G_v (resp. G_e).

A *generalized Baumslag-Solitar* (GBS) group is a group which is isomorphic to the fundamental group of a finite graph of groups where all vertex and edge groups are infinite cyclic. If a generator is chosen for every vertex and edge group then the monomorphisms ϕ_e are defined by the multiplication by an integer $\lambda(\bar{e}) \in \mathbb{Z} \setminus \{0\}$.

Let G be a GBS group. In the general case there exist infinitely many marked graphs of cyclic groups for G. Making $\mathrm{Aut}(G)$ act on a marked graph of groups often yields infinitely many other markings, besides in general cases there are infinitely many possible underlying graphs of groups.

Let T be a G-tree. A subgroup $H < G$ is *elliptic* in T if it fixes a point in T. Suppose all elliptic groups in T are also elliptic in S. Then there exists a G-equivariant map $T \to S$ (see by example [GL07]).

Let T, T' be G-trees. We say that they lie in the same *deformation space* if they have the same sets of elliptic subgroups. Equivalently they are in the same deformation space if there exist G-equivariant maps $T \to T'$ and $T' \to T$.

Now let us define the *cyclic deformation space* $\mathcal{D}$ associated to a group G, as the set of minimal simplicial G-trees with cyclic vertex and edge stabilizers, where we identify T and T' if there is a G-equivariant isometry or homothety $T \to T'$.

Equivalently we could define $\mathcal{D}$ with marked graphs of groups.

If G is not isomorphic to $\mathbb{Z}^2$ or the fundamental group of a Klein bottle $\langle a, b | a^2 = b^2 \rangle \simeq \langle a, t | tat^{-1} = a^{-1} \rangle$, then $\mathcal{D}$ is a deformation space, i.e. all trees in $\mathcal{D}$ have the same elliptic subgroups.

Let $T \in \mathcal{D}$ and let e be an edge in T. Define the equivalence relation $\sim_e$ as the minimal G-invariant equivalence relation such that $x \sim_e y$ if $x, y \in e$. The *collapse* of the edge e is the quotient map $T \to T/\sim_e$. The edge e in T is *collapsible* if $T/\sim_e \in \mathcal{D}$. Equivalently an edge is collapsible if its image in the quotient is not a loop and one of its two labels is ± 1.

A tree of $\mathcal{D}$ is *reduced* if none of its edges is collapsible.

GBS trees in the same deformation space share some properties. Let Γ be a finite connected graph. Then the first Betti number $b_1(\Gamma)$ is defined by $b_1(\Gamma) = \#E(\Gamma) - \#V(\Gamma) + 1$. By [GL07, Section 4] the first Betti number is an invariant of the deformation space.

Let G be a GBS group with cyclic deformation space $\mathcal{D}$. We say that an elliptic subgroup $H < G$ is *big* if there exists a tree $T \in \mathcal{D}$ such that H fixes no edge in T.

From [GL07] we deduce:

Lemma 3.1.1. *Let $T \in \mathcal{D}$. The number of vertices $v \in T$ such that for all edge e with origin v, $G_e \neq G_v$ is bounded by the number of conjugacy classes of big subgroups of G.*

If T is reduced then these numbers are equal.

Remark 3.1.2. The notion of big subgroups is defined in [GL07], though it depends on a family $\mathcal{A}$ of subgroups of G: a subgroup $H < G$ is big if it is elliptic and is not conjugate into a subgroup of an element of $\mathcal{A}$. Here the corresponding choice for $\mathcal{A}$ is the family of subgroups which fix an edge in a reduced tree of $\mathcal{D}$, or equivalently in every reduced tree of $\mathcal{D}$. Thus an elliptic subgroup is big if fixes a single point in some (equivalently any) tree in $\mathcal{D}$.

Solvable GBS groups are GBS groups isomorphic to $\mathbb{Z}$ and $\mathrm{BS}(1, n)$ for $n \in \mathbb{N}$ (which include $\mathbb{Z}^2$ and the fundamental group of a Klein bottle).

3.1.2 Cyclic factors, irreducible automorphisms

From now on we assume G is a non-solvable GBS group. The automorphism group of G is $\mathrm{Aut}(G)$. The outer automorphism group is $\mathrm{Out}(G) := \mathrm{Aut}(G)/\mathrm{Inn}(G)$ where $\mathrm{Inn}(G)$ is the subgroup of inner automorphisms $\{c_g : x \mapsto gxg^{-1}, g \in G\}$.

Cyclic factors are the GBS analogue of free factors for free groups.

Definition 3.1.3. A *cyclic factor* of G is a subgroup H such that there exists a graph of cyclic groups Γ and an identification $G \simeq \pi_1(\Gamma)$, with a subgraph Γ_H such that H is conjugate to $\pi_1(\Gamma_H)$.

The family of cyclic factors of G is stable by conjugacy and by automorphisms.

Examples 3.1.4.

1. If $G := \mathrm{BS}(2, 4) = \langle a, t | ta^2t^{-1} = a^4 \rangle$, the first graph of groups of Figure 3.1 represents G. The red subgraph represents the subgroup $H := \langle a, t^{-1}a^2t \rangle$ which is a cyclic factor.
2. The second and third graphs of Figure 3.1 represents $G := \langle u, r, s, t | tu^nt^{-1} = su^ns^{-1} = ru^nr^{-1} = u \rangle$. The subgroup $\langle u, r \rangle \simeq \mathrm{BS}(1, n)$ is a cyclic factor, it can be seen in the graph on the left. The subgroup $\langle u, rur^{-1}, rsus^{-1}r^{-1}, rst \rangle \simeq \mathrm{BS}(1, n^3)$ is a cyclic factor which can be seen in the graph on the right.

Definition 3.1.5. An automorphism $\phi \in \mathrm{Aut}(G)$ is *fully irreducible* if no power of ϕ preserves the conjugacy class of a cyclic factor. Since inner automorphisms preserve conjugacy classes, the full irreducibility can be defined for outer automorphisms.

Example 1. ($G \simeq BS(2,4)$)

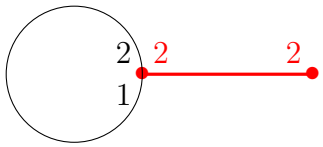

Example 2.

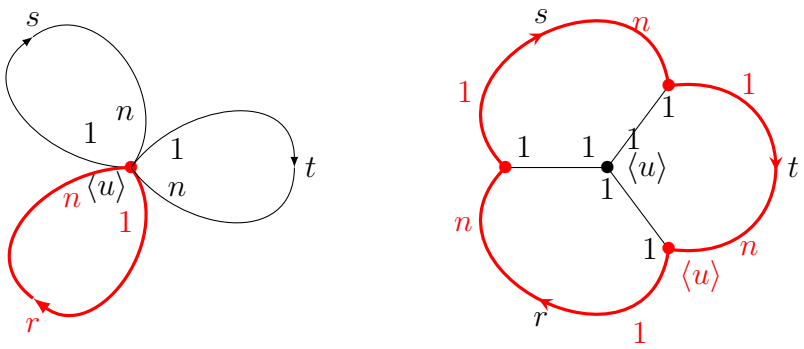

Figure 3.1 – Examples of cyclic factors

A *representative* for ϕ is a map $f : T \to T$ with $T \in \mathcal{D}$ which is ϕ-equivariant, i.e. $\forall t \in T, \forall g \in G,\ f(gt) = \phi(g) \cdot f(t)$.

A representative for an outer automorphism class $\psi \in \mathrm{Out}(G)$ is a representative for some automorphism in the class ψ.

Definition 3.1.6. Let $\phi \in \mathrm{Out}(G)$. A *pseudo-periodic* conjugacy class for ϕ is the conjugacy class of an element $g \in G$ such that $\|\phi^n(g)\|_{n\in\mathbb{N}}$ is bounded.

An automorphism $\phi \in \mathrm{Out}(G)$ is *pseudo-atoroidal* if ϕ has no pseudo-periodic conjugacy class.

Train track representatives for automorphisms of $\mathrm{Out}(F_N)$ were introduced in [BH92]. They are a main tool for studying these automorphisms. One can define likewise train tracks for other families of groups acting on trees.

Definition 3.1.7. Let T be a G-tree. A *gate structure* on T is a G-invariant family of equivalence relations $(\sim_v)_{v\in V(T)}$ on the sets E_v of edges with origin v. Equivalence classes for these relations are called gates.

Let $\tau : T \to T'$ be a G-equivariant map sending edges to non-degenerate non-backtracking paths. The gate structure *induced by* τ is the minimal gate structure such that for $v \in V(T)$, $e, e' \in E_v$, if $\tau(e) \cap \tau(e')$ has non-zero length then $e \sim_v e'$.

Definition 3.1.8. Let T be a G-tree with a gate structure. A *turn* in T is a pair of edges with same origin. The turn $\{e, e'\}$ is *illegal* if e and e' belong to the same gate. Otherwise the turn is *legal*.

Definition 3.1.9. Let $\phi \in \mathrm{Out}(G)$. Let $f : T \to T$ be a representative for ϕ sending vertex to vertex.. Then f is *train track* if for every $e \in E(T)$:

— $\mathrm{len}(f(e)) > 0$

— at every vertex $v \in V(T)$ there are at least two gates for the gate structure induced by f.

— for every $k \in \mathbb{N}$, $f^k(e)$ crosses only legal turns for the gate structure induced by f.

The *train track structure* is the gate structure such that $e \sim e'$ if there exists $k \in \mathbb{N}$ such that $f^k(e) \cap f^k(e')$ is not a single point.

Suppose $f : T \to T$ is a train track representative for a fully irreducible automorphism $\phi \in \mathrm{Out}(G)$. Up to precomposing f with a map $T \to T$ whose restriction to edges is a homeomorphism, we may assume that f stretches the edges uniformly. Let $e_1, \ldots, e_n$ be the edges of T/G. Let $A(f)$ be the transition matrix for f where $A(f)_{ij}$ is the number of occurences of edges in the orbit e_i in $f(e_j)$, for $1 \leq i, j \leq n$. Irreducibility of ϕ implies that up to collapsing edges in T the matrix $A(f)$ is primitive, i.e. there exists $k \in \mathbb{N}$ such that $A(f)^k > 0$ (see Chapter 2, Lemma 2.1.9).

The theorem of Perron-Frobenius below then applies to $A(f)$:

Theorem 3.1.10 (Perron-Frobenius)**.** *Let A be a non-negative primitive matrix with size $n \times n$. There exists a real eigenvalue $\lambda > 0$ (the* Perron-Frobenius eigenvalue*) such that for every other eigenvalue $\mu \neq \lambda$ we have $|\mu| < \lambda$. The eigenvectors for λ are unique up to scalar multiplication and there exists an eigenvector v for λ such that $v > 0$.*

A proof of the theorem can be found in [Sen81, Theorem 1.1].

Let λ be the Perron-Frobenius eigenvalue and $(l_1, \ldots, l_n)$ be the left Perron-Frobenius eigenvector. Define a metric on T by $\mathrm{len}(e_n) := l_n$. Then for every $e \in E(T)$ we have $\mathrm{len}(f(e)) = \lambda\,\mathrm{len}(e)$, and the Lipschitz constant $\mathrm{Lip}(f)$ is λ.

From now on we assume train track maps are linear on edges and have the same Lipschitz constant on all edges.

Recall that it is not known whether all fully irreducible automorphisms admit train track maps. However some do exist: Example 3.1.11 is an example of fully irreducible automorphism with a train track representative.

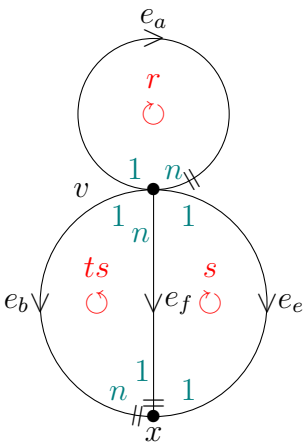

Figure 3.2 – The quotient T/G in Example 3.1.11

Examples 3.1.11. Let $G := \langle r, s, t | ru^n r^{-1} = su^n s^{-1} = atu^n t^{-1} = u \rangle$. Define $\phi \in \mathrm{Aut}(G)$ by

$$\phi : \begin{cases} u \mapsto u \\ r \mapsto s \\ s \mapsto t \\ t \mapsto rsts^{-1}t^{-1} \end{cases}$$

Define the tree T (whose quotient T/G is represented on Figure 3.2) by a fundamental domain with vertices v, x and edges

$$\begin{aligned} e_a &= [v, rv] \\ e_b &= [v, tx] \\ e_e &= [v, x] \\ e_f &= [v, s^{-1}x] \end{aligned}$$

Define f on T by

$$
\begin{aligned}
f :& v \mapsto v, x \mapsto tx \\
& e_a \mapsto e_e - s \cdot \bar{e}_f \\
& e_b \mapsto e_a - r \cdot e_e - rs \cdot \bar{e}_f - rs \cdot e_b - rst \cdot \bar{e}_e - rst \cdot e_f \\
& e_e \mapsto e_b \\
& e_f \mapsto e_e
\end{aligned}
$$

The map f is a train track representative for ϕ. To see this we can compute the successive images of the turns taken by f. Consider the turn $\{\bar{e}_a, re_e\}$ at vertex rv which is taken by $f(e_b)$:

$$
r\{r^{-1}\bar{e}_a, e_e\} \xrightarrow{f} \phi(r)\{e_f, e_b\} \xrightarrow{f} \phi^2(r)\{e_e, e_a\} \xrightarrow{f} \phi^3(r)\{e_b, e_e\} \xrightarrow{f} \phi^4(r)\{e_a, e_b\}
$$

$$
\phi^4(r)\{e_a, e_b\} \xrightarrow{f} \phi^2(r)\{e_e, e_a\}
$$

When applying f to $\phi^4(r)\{e_a, e_b\}$ we find a turn in the orbit of $\phi^2(r)\{e_e, e_a\}$ which we had already found. Thus $\{\bar{e}_a, re_e\}$ is a legal turn. The proof for the other turns goes the same way.

The inverse of ϕ is

$$
\phi^{-1} : \begin{cases} u \mapsto u \\ r \mapsto tsrs^{-1}r^{-1} \\ s \mapsto r \\ t \mapsto s \end{cases}
$$

which is the same automorphism as ϕ, but with swapped roles for r and t. Thus it admits a train track map.

Let $f : T \to S$ be a G-equivariant map between trees of $\mathcal{D}$, sending vertex to vertex and edge to edge path. By [For02] it is a quasi-isometry, i.e. a map such that for all $x, y \in T$, $K^{-1}d_T(x, y) - C \leq d_{T'}(h(x), h(y)) \leq Kd_T(x, y) + C$. Thus the map has bounded backtracking property (see [Gab+98]) and there exists a constant $\mathrm{BBT}(f) \leq K^2C + C$ such that for every $x, y \in T$ the image $f([x, y])$ lies in a $\mathrm{BBT}(f)$-neighbourhood of $[f(x), f(y)]$.

In [BFH97] and [Alg11], a similar constant $\mathrm{BCC}(f)$ is used. Let us clarify the link between these constants. Suppose $\alpha \cdot \beta$ is a geodesic concatenation of paths with α f-legal. When applying f to $\alpha \cdot \beta$, there is a subsegment $\tau \subset \alpha$ such that $f(\tau) \subset f(\beta)$. Let

$\tau' \subset \beta$ be a minimal prefix such that $f(\tau) \subset f(\tau')$. The first point of τ and last point of τ' are mapped to the same point. Then because f is a quasi-isometry there is a constant K depending only on f such that the length of $\tau \cdot \tau'$ is bounded by K. This gives a bound $\mathrm{BCC}(f)$ (from *bounded cancellation constant*, introduced by Cooper in [Coo87]) on $f(\tau)$: the simplification which occurs at an illegal turn is bounded by $\mathrm{BCC}(f)$. By applying the previous paragraph to $\tau \cdot \tau'$ we have $\mathrm{BCC}(f) \leq \mathrm{BBT}(f)$. In order to keep a reduced number of constants, we will use the BBT constant instead of BCC where it could be used.

For a path η in a tree T, we denote by $[\eta]$ the unique non-backtracking path which has the same endpoints as η. We can extend this notation to infinite paths which converge to a point in the boundary of T: for example, this is well-defined for a bi-infinite quasi-geodesic such as the image of a line by a quasi-isometry.

Definition 3.1.12. Let $\phi \in \mathrm{Out}(G)$ be a fully irreducible automorphism with a train track representative $f : T \to T$. The *critical constant* is $C_f := \frac{2\,\mathrm{BBT}(f)}{\lambda - 1}$.

The critical constant has the following property: for any geodesic concatenation $\alpha \cdot \beta \cdot \gamma \subset T$ such that β is legal and $\mathrm{len}(\beta) \geq C_f$ then let $\alpha' \subset \alpha, \beta' \subset \beta, \gamma' \subset \gamma$ such that the path $[f(\alpha \cdot \beta \cdot \gamma)]$ can be written as the concatenation $[f(\alpha')] \cdot [f(\beta')] \cdot [f(\gamma')]$. Then $\mathrm{len}(f(\beta')) \geq C_f$. More specifically we have:

Lemma 3.1.13. *Let $\phi \in \mathrm{Out}(G)$ be fully irreducible. Let $f : T \to T$ be a train track representative. Let α be any path in T. Let $\beta \subset \alpha$ be a legal subpath with length at least $2C_f$. Define β' as the legal subpath of β obtained by cutting out the $\frac{C_f}{2}$-neighbourhood of the endpoints. Then β' satisfies the following condition: for all $n \in \mathbb{N}$*

$$f^n(\alpha \setminus \beta') \cap f^n(\beta') = \varnothing \text{ and } f^n(\beta') \subset [f^n(\alpha)]$$

In particular $\mathrm{len}([f^n(\alpha)]) \geq \lambda^n \,\mathrm{len}(\beta')$.

Proof. Let α, β be as above.

We actually prove a slightly stronger statement by induction on n. Define β'_n as the segment obtained by cutting out a $\mathrm{BBT}(f) \sum_{k=1}^{n} \lambda^{-k}$-neighbourhood from the endpoints of β. Observe that $\beta' = \bigcap_{n \in \mathbb{N}} \beta'_n$ since $\mathrm{BBT}(f) \sum_{k=1}^{\infty} \lambda^{-k} = \dfrac{\mathrm{BBT}(f)}{\lambda - 1} = \dfrac{C_f}{2}$.

We will prove that for all $n \in \mathbb{N}$ we have

$$f^n(\alpha \setminus \beta'_n) \cap f^n(\beta'_n) = \varnothing \text{ and } f^n(\beta'_n) \subset [f^n(\alpha)]$$

and since $\beta' \subset \beta'_n$ is a subsegment of the legal segment β'_n, the same holds for β'.

The condition is true for $n = 0$ since $\beta'_0 = \beta$.

Suppose the lemma holds for some $n \in \mathbb{N}$. Since β'_n is legal we have $\operatorname{len}(f^n(\beta'_n)) = \lambda^n \operatorname{len}(\beta'_n)$. Besides $f^n(\beta'_n) \subset [f^n(\alpha)]$. Apply f to the path $f^n(\alpha)$. There may be cancellation at the endpoints of the legal segment $f^n(\beta'_n)$ but this cancellation does not exceed $\mathrm{BBT}(f)$ when measured in $f^{n+1}(\beta'_n)$, since β'_n is legal.

This neighbourhood in $f^{n+1}(\beta'_n)$ corresponds to a $\lambda^{-n-1}\,\mathrm{BBT}(f)$-neighbourhood of the endpoints of β'_n: as a result $f^{n+1}(\beta'_{n+1})$ does not intersect $f(f^n(\alpha) \setminus f^n(\beta'_n))$ since it is contained in β'_n. It does not intersect $f(f^n(\beta'_n) \setminus f^n(\beta'_{n+1}))$ either. Finally note that $f(f^n(\alpha) \setminus f^n(\beta'_n)) \cap f^{n+1}(\beta'_n) = f^{n+1}(\alpha \setminus \beta'_n) \cap f^{n+1}(\beta'_n)$.

Since $f^{n+1}(\beta'_{n+1})$ does not meet $f^{n+1}(\alpha \setminus \beta'_{n+1})$, is is contained in $[f^{n+1}(\alpha)]$. □

A consequence of Lemma 3.1.13 is:

Lemma 3.1.14. *Let $\phi \in \mathrm{Out}(G)$ be fully irreducible. Let $f : T \to T$ be a train track representative. For any path α in T, for any legal subpath $\beta \subset \alpha$ such that $\operatorname{len}(\beta) > 2C_f$, we have for every $n \in \mathbb{N}$:*

$$\operatorname{len}(f^n(\beta) \cap [f^n(\alpha)]) \geq \frac{1}{2}\lambda^n \operatorname{len}(\beta)$$

Proof. Suppose α contains a legal subsegment β of length greater than $2C_f$. Let β' be the subsegment of β at distance $\frac{C_f}{2}$ from the endpoints of β. By Lemma 3.1.13, for every $n \in \mathbb{N}$, $f^n(\beta') \subset [f^n(\alpha)]$ so we have:

$$\begin{aligned}
\operatorname{len}(f^n(\beta')) &= \lambda^n \operatorname{len}(\beta') \\
&= \lambda^n \left(\operatorname{len}(\beta) - \frac{2\,\mathrm{BBT}(f)}{\lambda - 1} \right) \\
&\geq \lambda^n \operatorname{len}(\beta) \left(1 - \frac{1}{2} \right)
\end{aligned}$$

Therefore $\operatorname{len}(f^n(\beta) \cap [f^n(\alpha)]) \geq \frac{1}{2}\lambda^n \operatorname{len}(\beta)$. □

Definition 3.1.15. Let $f : T \to T$ be a train track representative for $\phi \in \mathrm{Out}(G)$. A non-backtracking segment $\eta \subset T$ is a *periodic Nielsen path* if there exists $g \in G$ and $n \geq 1$ such that $g[f^n(\eta)] = \eta$. We call it simply a Nielsen path if we can choose $n = 1$.

A periodic Nielsen path is *indivisible* if it cannot be written as the non-backtracking concatenation $\alpha \cdot \beta$ of two periodic Nielsen paths.

A result about periodic indivisible Nielsen paths (or *pINPs*) from Chapter 2 is:

Proposition 3.1.16. *Let $f : T \to T$ be a train track representative for an automorphism $\phi \in \mathrm{Out}(G)$. There are only finitely many orbits of periodic indivisible Nielsen paths.*

Periodic Nielsen paths give a characterization of pseudo-periodic conjugacy classes for $\phi \in \mathrm{Out}(G)$, proved in Chapter 2, Section 2.3:

Lemma 3.1.17. *Let $f : T \to T$ be a train track representative for $\phi \in \mathrm{Out}(G)$. The conjugacy class of an element $g \in G$ is pseudo-periodic for ϕ if and only if the axis of g in T is a geodesic concatenation of periodic indivisible Nielsen paths.*

Lemma 3.1.18. *Suppose that ϕ is fully irreducible, pseudo-atoroidal and has a train track representative $f : T \to T$. Then there exists $m \in \mathbb{N}$ such that it is impossible to concatenate more than m periodic Nielsen paths for f together in T, and more than m periodic Nielsen paths for f_- in T_-.*

Proof. There are finitely many orbits of periodic INPs; let l be the number of orbits of periodic INPs. By contradiction, suppose L is a path in T which contains a concatenation of more than $2l$ pINPs. Then there exists $\eta, g\eta \subset L$ with $g \in G$ loxodromic such that there is a fundamental domain for g which is a concatenation of pINPs. By Lemma 3.1.17, this implies that g is a loxodromic pseudo-periodic element for ϕ, which is impossible since ϕ is pseudo-atoroidal, so we can set $m := 2l$. □

3.1.3 The Lipschitz metric on $\mathcal{D}$

The space $\mathcal{D}$ can be endowed with a pseudo-metric called the *Lipschitz metric*. For $T, T' \in \mathcal{D}$ define

$$\mathrm{Lip}(T, T') = \inf_{f:T\to T'} \mathrm{Lip}(f)$$

where the infimum is taken over all G-equivariant Lipschitz functions $f : T \to T'$. In [Mei15] the following is proved:

Proposition 3.1.19. *For $T, T' \in \mathcal{D}$ there exists a G-equivariant map $f : T \to T'$, sending vertex to vertex and edge to edge path, linear on the edges, such that* $\mathrm{Lip}(f) = \mathrm{Lip}(T, T')$.

The Lipschitz metric is defined as follows: for $T, T' \in \mathcal{D}$

$$d_{\mathrm{Lip}}(T, T') = \log\left[\mathrm{Lip}(T, T')\frac{\mathrm{vol}(T/G)}{\mathrm{vol}(T'/G)}\right]$$

The distance $d_{\mathrm{Lip}}(T, T')$ is unchanged by rescaling T or T': it only depends on their projective classes. If T, T' are normalized so that $\mathrm{vol}(T) = \mathrm{vol}(T') = 1$ then $d_{\mathrm{Lip}}(T, T') = \log \mathrm{Lip}(T, T')$. Sometimes it is more practical to work with 1-Lipschitz maps, for example when $T \to T'$ is a collapse or a fold. When $\mathrm{Lip}(T, T') = 1$ then $d_{\mathrm{Lip}}(T, T') = \log \frac{\mathrm{vol}(T)}{\mathrm{vol}(T')}$.

The Lipschitz metric is not a metric in the actual sense.

Lemma 3.1.20. *The Lipschitz metric has the following properties:*

(i) for $T, T' \in \mathcal{D}$, $d_{\mathrm{Lip}}(T, T') \geq 0$

(ii) for $T, T', T'' \in \mathcal{D}$, $d_{\mathrm{Lip}}(T, T'') \leq d_{\mathrm{Lip}}(T, T') + d_{\mathrm{Lip}}(T', T'')$

Proof. (i) Let $T, T' \in \mathcal{D}$ scaled such that $\mathrm{Lip}(T, T') = 1$. Then f induces a 1-Lipschitz map on the quotients. By minimality f is surjective so $\mathrm{vol}(T'/G) \leq \mathrm{vol}(T/G)$. Thus

$$\log \mathrm{Lip}(T, T') \frac{\mathrm{vol}(T/G)}{\mathrm{vol}(T'/G)} \geq 1$$

(ii) Let $T, T', T'' \in \mathcal{D}$. Let $f : T \to T', f' : T' \to T''$ be Lipschitz maps. Then $\mathrm{Lip}(T, T'') \leq \mathrm{Lip}(f' \circ f) \leq \mathrm{Lip}(f)\,\mathrm{Lip}(f')$. By taking the lower bound we get $\mathrm{Lip}(T, T'') \leq \mathrm{Lip}(T, T')\,\mathrm{Lip}(T', T'')$. By taking the logarithm we obtain what we want. □

Remark 3.1.21. The other properties of metrics fail for $\mathcal{D}$:

— like in CV_N the Lipschitz metric is not symmetric. A common counter example is drawn on Figure 3.3: if T, T' are the same tree with a different metric on edges such that T'/G has a very short loop, $d_{\mathrm{Lip}}(T', T)$ is very big.

— unlike in CV_N there exist T, T' in $\mathcal{D}$ such that $d_{\mathrm{Lip}}(T, T') = 0$ and $d_{\mathrm{Lip}}(T', T) \neq 0$. More precisely, $d_{\mathrm{Lip}}(T', T)$ can be chosen arbitrarily big. See Figure 3.3.

— If $\phi \in \mathrm{Out}(G)$ and $T \in \mathcal{D}$ then $d_{\mathrm{Lip}}(T, T \cdot \phi) = 0$ implies that $T = T \cdot \phi$. This is a consequence of the fact that the actions are minimal so maps are surjective. A 1-Lipschitz G-equivariant surjective map from T to itself is an isometry, unless G is solvable.

The Lipschitz metric and its computation have been explored before in [Bes11] for free groups, [FM15] for free products, and [Mei15] for more general deformation spaces. The facts presented below can be found in these papers.

Let $T, T' \in \mathcal{D}$. Let $f : T \to T'$ be a piecewise linear G-equivariant map. The *tension graph* $\Delta(f)$ is the subforest of T spanned by edges $e \in E(T)$ such that the stretch factor

1. Classical example for F_2

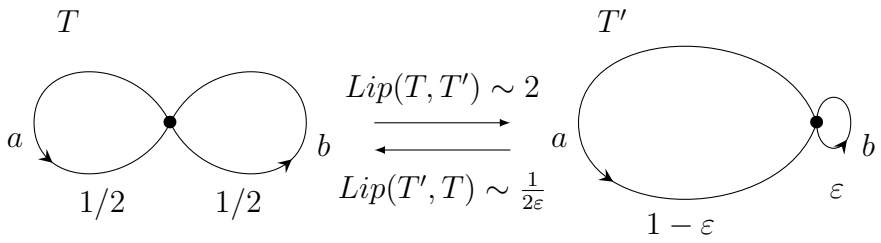

2. Example in $BS(2,4)$

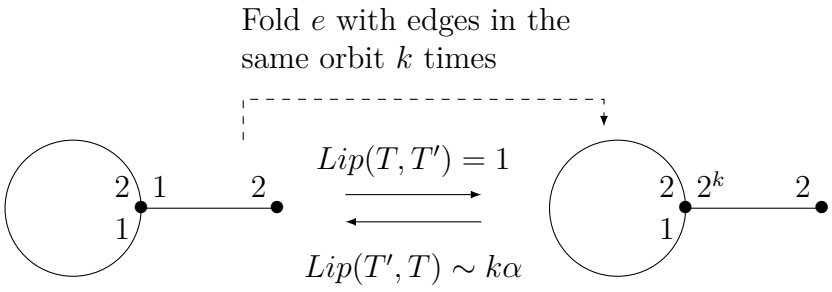

Figure 3.3 – Counter examples for symmetry of the Lipschitz metric

on e is $\mathrm{Lip}(f)$. The map f is *optimal* if it realizes the infimum of $\mathrm{Lip}(T, T')$ and if at every vertex $v \in \Delta(f)$, there are at least two gates at v for the gate structure induced by f which contain edges in $\Delta(f)$. In [Mei15] Meinert proves that optimal maps exist.

The distance between two points in $\mathcal{D}$ can be effectively computed by comparing translation lengths of some elements of G in both trees. Suppose $f : T \to T'$ is a G-equivariant map between trees of $\mathcal{D}$. For every $g \in G$ we have, by applying f to a fundamental domain,

$$\frac{\|g\|_{T'}}{\|g\|_T} \leq \mathrm{Lip}(f)$$

thus by taking the lower bound we have $\mathrm{Lip}(T, T') \geq \max_{g \in G} \frac{\|g\|_{T'}}{\|g\|_T}$. It is actually an equality, as this result from [Mei15] states:

Lemma 3.1.22. *Let $T, T' \in \mathcal{D}$. Let $f : T \to T'$ be an optimal map. There exists $g \in G$ such that $\mathrm{Axis}_T(g)$ is f-legal and contained in the tension graph for f. In particular*

$$\mathrm{Lip}(T, T') = \frac{\|g\|'_T}{\|g\|_T} = \max_{h \in G} \frac{\|h\|'_T}{\|h\|_T}$$

Let $T \in \mathcal{D}$. A *candidate* of T is an element $g \in G$ such that the map $\pi : \mathrm{Axis}_T(g)/\langle g \rangle \to T/G$ has one of the following forms (see Figure 3.4):

— a loop: the map π is an embedding

— a figure eight: there are two embedded circles u, v in T which intersect in exactly one point. The map π maps the circle $\mathrm{Axis}_T(g)/\langle g \rangle$ to the tight loop which crosses u and v successfully.

— a barbell: there are two disjoint embedded circles u, v in T/G, and a segment s which connects u to v; π maps the circle to the tight loop which crosses u, s, then v, and then s backwards.

— an embedded singly degenerate barbell : it is the degenerate case of the barbell where v is a single vertex. In that case, the vertex group at v must be greater than the edge group of the last edge of s.

— an embedded doubly degenerate barbell : degenerate case of the barbell where both circles are single vertices. The vertex group at u must also be greater than the edge group of the first edge of g.

In particular a candidate crosses each orbit of edge at most twice. In [FM15, Theorem 9.10] the following theorem is proved for the case of free products:

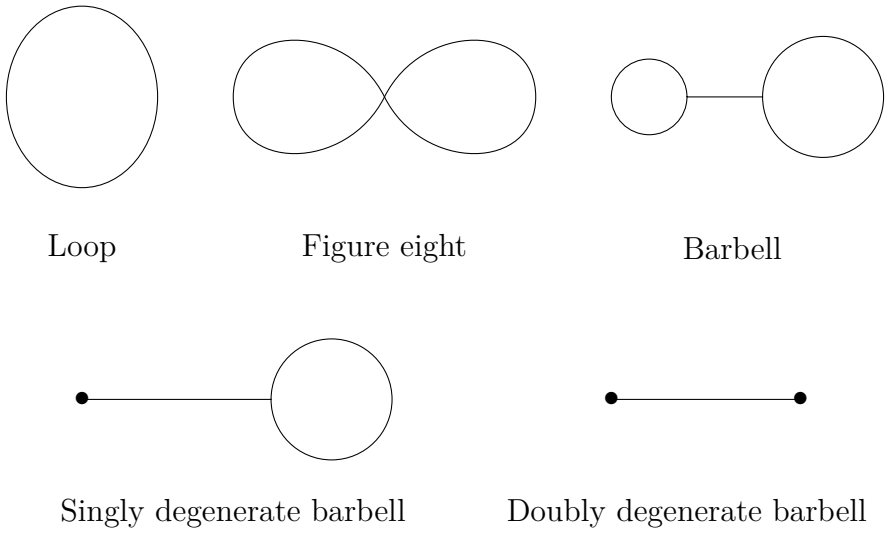

Figure 3.4 – The five possible shapes for a candidate in the quotient graph

Theorem 3.1.23. *Let $T \in \mathcal{D}$. For every $T' \in \mathcal{D}$, there exists a candidate g of T such that*

$$\mathrm{Lip}(T, T') = \frac{\|g\|'_T}{\|g\|_T}$$

For the proof we refer to [FM15]. Note that the context differs a little since the deformation space and the group are different. The relevant point is that contrary to elements of CV_N, trees in the outer space for a free product may have non-free vertex stabilizers, which account for the degenerate barbells. In the case of GBS products we also have vertices whose stabilizer is greater than the stabilizers of incident edges, hence the presence of degenerate barbell candidates.

3.1.4 The axis of an irreducible automorphism

Proposition 3.1.24. *For an automorphism $\phi \in \mathrm{Out}(G)$ with a primitive train track map $f : T \to T$, then*

$$\mathrm{Lip}(f) = \mathrm{Lip}(T, T \cdot \phi) = \min_{S \in \mathcal{D}} \mathrm{Lip}(S, S \cdot \phi)$$

Proof. Since f is a train track map, at each vertex of T there are at least two gates. Consequently there exists $h \in G$ such that $\mathrm{Axis}_T(g)$ is legal. Let $n \in \mathbb{N}$. Since h is f^n-legal for all $n \in \mathbb{N}$ we have $\frac{\|\phi^n(h)\|_T}{\|h\|_T} = \mathrm{Lip}(f^n) = \mathrm{Lip}(f)^n$.

Then we have $\mathrm{Lip}(T, T \cdot \phi^n) = \sup_{g \in G} \frac{\|\phi^n(g)\|_T}{\|g\|_T} = \frac{\|\phi^n(h)\|_T}{\|h\|_T} = \mathrm{Lip}(f)^n$. Let $\lambda := \mathrm{Lip}(f)$.

Let $S \in \mathcal{D}$. Let $f' : S \to S \cdot \phi$ be such that $\mathrm{Lip}(S, S \cdot \phi) = \mathrm{Lip}(f') =: \lambda'$. By triangular inequality, for $n \in \mathbb{N}$ we have $\mathrm{Lip}(S, S \cdot \phi^n) \leq \lambda'^n$.

By triangular inequality we have

$$\lambda^n = \mathrm{Lip}(T, T \cdot \phi^n) \leq \mathrm{Lip}(T, S)\,\mathrm{Lip}(S, S \cdot \phi^n)\,\mathrm{Lip}(S, T) \leq \lambda'^n\,\mathrm{Lip}(T, S)\,\mathrm{Lip}(S, T)$$

If $\lambda' < \lambda$ this inequality becomes false when n is big enough, hence the minimality of $\lambda = \mathrm{Lip}(T, T \cdot \phi)$. □

A map $\gamma : \mathbb{R} \to \mathcal{D}$ is a *geodesic* for the Lipschitz distance if for any $t, s \in \mathbb{R}$, $t < s \Rightarrow d_{\mathrm{Lip}}(\gamma(t), \gamma(s)) = s - t$. Since the metric is not symmetric, the distance $d_{\mathrm{Lip}}(\gamma(s), \gamma(t))$ needs not be $|s - t|$: in fact it can even be zero.

Proposition 3.1.25. *Let $f : T \to T$ be a train track representative for ϕ. There exists a geodesic $\mathcal{L}_f : \mathbb{R} \to \mathcal{D}$ such that for every $n \in \mathbb{Z}$, $T \cdot \phi^n \in \mathcal{L}_f$, and such that the map $\mathcal{L}_f$ is continuous for the axes topology.*

For a construction see [Mei15, Section 4.4]. The *axes topology* on $\mathcal{D}$ is the coarsest topology such that the functions $T \mapsto \|g\|_T$ for $g \in G$ are continuous; for more information on the topologies of $\mathcal{D}$ see [GL07].

If $f : T \to T$ is a train track representative we choose an arbitrary axis $\mathcal{L}_f$ crossing T and denote by T_t the unique point of the axis such that $d_{\mathrm{Lip}}(T, T_t) = t$ if $t \geq 0$ and $d_{\mathrm{Lip}}(T_t, T) = -t$ if $t \leq 0$.

Similarly there is an axis $\mathcal{L}_{f_-}$ for ϕ^{-1} defined from $f_- : T_- \to T_-$. Since d_{Lip} is $\mathrm{Out}(G)$-invariant the axes stay within a bidirectional bounded neighbourhood of each other.

We already stated that between any two trees in $\mathcal{D}$, there exists a G-equivariant quasi-isometry. The quasi-isometry constants can be chosen uniformly if the trees lie in a subsegment of $\mathcal{L}_f$:

Lemma 3.1.26. *For any $T, S \in \mathcal{D}$ and $D \geq 1$ such that $\max\{\mathrm{Lip}(T, S), \mathrm{Lip}(S, T)\} \leq D$, any optimal map $T \to S$ is an equivariant $(D, 4D)$-quasi-isometry.*

Proof. Let $T, S \in \mathcal{D}$. Scale T, S such that $\mathrm{vol}(T/G) = \mathrm{vol}(S/G) = 1$. There exists $X \subset T$ such that $G \cdot X = T$ and $\mathrm{diam}(X) \leq 1$. For every oriented edges $\overrightarrow{e}, \overrightarrow{e'} \in E(T)$, there exists an edge path with length at most 2 with first edge $\overrightarrow{e}$ and last edge $h\overrightarrow{e'}$ for some $h \in G$. This fails if G is solvable, but we assumed that it is not the case.

Let us prove this fact. First we will prove that for any edge e there exists a path $\overrightarrow{e} \dots \overleftarrow{he}$ for some $h \in G$.

Since G is not solvable, then the action of G on T is irreducible. Since T is not a line, there exist translates $h_1 e, h_2 e$ such that for any line containing both edges, the orientations of the edges along the line differ.

There exists $h \in G$ such that $\mathrm{Axis}_T(h)$ crosses both $h_1 e$ and $h_2 e$. Either $h_1 e$ and $h_2 e$ point towards each other, or they may point away from each other, but in that case there exists $k \in \mathbb{Z}$ such that $h^k h_1 e$ and $h_2 e$ point towards each other. Then we get a path $\overrightarrow{h_1 e} \dots \overleftarrow{h_2 e}$.

Now let e, e' be edges in the tree. There is a path which connects both edges, but the path may fail to contain a translate of $\overrightarrow{e} \dots \overrightarrow{he'}$ for some $h \in G$. By concatenating paths which reverse the orientation on one or both sides we obtain a path satisfying the condition.

As for the bound on the length of the path, observe that if an edge appears in the path twice with same orientation, then subpath can be deleted to obtain a shorter path.

Let $\tau : T \to S$ be an optimal map. Let $x, y \in T$. Let e be the first edge of $[x, y]$. Let e' be an edge with origin y and not in $[x, y]$: such an edge exists since T is minimal. By the fact above, there exists $g \in G$ and a path containing $\overrightarrow{e}, g\overrightarrow{e}$ with these orientations and length at most 2. Thus $d_T(gx, y) \leq 2$ and $[x, y] \subset \mathrm{Axis}_T(g)$. We have $d_T(x, y) + 2 \leq \|g\|_T$.

Then by Lemma 3.1.27 we have $\mathrm{Lip}(S, T) \leq D, \mathrm{Lip}(T, S) \leq D$ so

$$d_S(\tau(x), \tau(y)) \leq D d_T(x, y)$$

and

$$\begin{aligned} d_S(\tau(x), \tau(y)) &\geq d_S(\tau(x), \tau(gx)) - d_S(\tau(gx), \tau(y)) \\ &\geq \|g\|_S - 2\,\mathrm{Lip}(\tau) \\ &\geq \frac{\|g\|_T}{D} - 2\,\mathrm{Lip}(\tau) \\ &\geq \frac{d_T(x, y)}{D} - \frac{2}{D} - 2D \end{aligned}$$

Then the optimal map $\tau : T \to T_{t'}$ is a $(D, 4D)$-quasi-isometry. □

We have the following result about the axes $\mathcal{L}_f$ and $\mathcal{L}_{f_-}$:

Lemma 3.1.27. *Let $\phi \in \mathrm{Out}(G)$ be a fully irreducible automorphism. Let $f : T \to T$ be a train track representative for ϕ and $f_- : T_- \to T_-$ be a train track representative for ϕ^{-1}. Let $\mathcal{L}_f, \mathcal{L}_{f_-}$ be axes in $\mathcal{D}$ for ϕ and ϕ^{-1}.*

Let $a < b$, c, d be real numbers. There exists a constant $D_{a,b,c,d} > 1$ such that for every $X, Y \in \{T_t/a \leq t \leq b\} \cup \{(T_-)_t/c \leq t \leq d\}$, $\mathrm{Lip}(X, Y) \leq D_{c,d}$.

Proof. Let $n, m \in \mathbb{N}$ be such that

$$n \log(\lambda) \leq a, m \log(\lambda) \geq b, n \log(\lambda_-) \leq c, m \log(\lambda_-) \geq d$$

There is a quadrilateron which crosses $T \cdot \phi^n, T \cdot \phi^m, T_- \cdot \phi^m, T_- \cdot \phi^n$, which contains $\{T_t/a \leq t \leq b\}$ and $\{(T_-)_t/c \leq t \leq d\}$. Its length is

$$d := d_{\mathrm{Lip}}(T \cdot \phi^n, T \cdot \phi^m) + d_{\mathrm{Lip}}(T, T_-) + d_{\mathrm{Lip}}(T_- \cdot \phi^m, T_- \cdot \phi^n) + d_{\mathrm{Lip}}(T_-, T)$$

Therefore, for every X, Y as above we have $L(X, Y) \leq D := e^d$. □

Remark 3.1.28. Lemmas 3.1.26 and 3.1.27 imply that for every $S \in \mathcal{D}$, there exist (K, C) such that for every $t \in [0, \log(\lambda)]$, there exist equivariant (K, C)-quasi-isometries $T_t \to S$ and $S \to T_t$.

3.2 The stable and unstable laminations

Let $T, T' \in \mathcal{D}$. There exists a G-invariant quasi-isometry $T \to T'$. In fact all equivariant quasi-isometries $T \to T'$ are close:

Lemma 3.2.1. *Let T, T' be metric G-trees such that T is co-compact. Let u, v be continuous G-equivariant maps $T \to T'$. There exists a constant C depending on u and v such that for every $x \in T$*

$$d_{T'}(u(x), v(x)) \leq C$$

Proof. Let $K \subset T$ be a compact subset such that $G \cdot K = T$. Let $C := \max_{x \in K} d_{T'}(u(x), v(x))$. For every $y \in T$ there exists $g \in G$ and $x \in K$ such that $y = gx$ so by G-equivariance

$$d_{T'}(u(y), v(y)) = d_{T'}(gu(x), gv(x)) = d_{T'}(u(x), v(x)) \leq C.$$

□

Recall that a G-invariant quasi-isometry f induces a G-equivariant homeomorphism $\partial T \to \partial T'$. Because of Lemma 3.2.1 the homeomorphism does not depend on f so there is a canonical G-invariant identification of the boundaries of all trees of $\mathcal{D}$.

A *lamination* Λ is a G-invariant, symmetric, closed subset of $\partial T \times \partial T \setminus \Delta$ where Δ is the diagonal, for some $T \in \mathcal{D}$. The discussion above implies that for any $S \in \mathcal{D}$, Λ can be canonically identified with a subset of $\partial S \times \partial S \setminus \Delta$ so we may drop the reference to T.

When we fix a tree T, Λ identifies with a G-invariant set of unoriented bi-infinite geodesics of T which we call the *realization* of Λ in T and which we denote by Λ_T. Its elements are called *leaves*. A *leaf segment* is a subsegment of a leaf of Λ_T. The assumption that Λ is a closed subset of $\partial T \times \partial T \setminus \Delta$ translates into the following fact: if $(\sigma_n)_{n\in\mathbb{N}}$ is an increasing sequence of leaf segments in T whose union is a bi-infinite geodesic $\ell \subset T$, then ℓ is a leaf of Λ_T.

If $T' \in \mathcal{D}$ is another tree, there exists a G-invariant quasi-isometry $h : T \to T'$. For any leaf $\lambda \in \Lambda_T$, the line of T' obtained by tightening $h(\lambda)$ is a leaf of $\Lambda_{T'}$, and conversely all leaves of $\Lambda_{T'}$ are tightened images of leaves of Λ_T.

Let $\phi \in \mathrm{Out}(G)$ be a fully irreducible automorphism. Let $f : T \to T$ be an primitive train track representative for ϕ with Lipschitz constant $\lambda > 1$. Define the *stable lamination* Λ_f^+ by its realization in the train track tree T, as the set of bi-infinite geodesics whose subsegments belong to

$$\{\sigma \subset T / \exists e \in E(T), \exists n \in \mathbb{N},\ \sigma \subset f^n(e)\}$$

We call $\Lambda_{f_-}^+$ the *unstable lamination*.

Remark 3.2.2. Since $\lambda > 1$, if e is an edge of T such that $e \subset f(\mathring{e})$, then the limit of $f^n(e)$ when $n \to \infty$ is a leaf of the lamination. Since f is primitive, one can check that the set of leaves which can be obtained by this process by replacing f with gf for $g \in G$ is a G-invariant subset of the stable lamination and its closure is the stable lamination.

Lemma 3.2.3. *For every $l > 0$ there exists $n_\Lambda > 0$ such that if $\alpha \subset T$ contains a legal subsegment with length at least $2C_f$, then $[f^n(\alpha)]$ contains a leaf segment of Λ_f^+ with length at least l for all $n \geq n_\Lambda$.*

Proof. Let β be a legal subsegment of α with length at least $2C_f$. By Lemma 3.1.13 the subpath $\theta \subset \beta$ obtained by truncating the $C_f/2$-neighbourhood of the endpoints has the following property: for any $n \in \mathbb{N}$, $f^n(\theta) \subset [f^n(\alpha)]$.

We have $\mathrm{len}(\theta) \geq C_f$. There exists $n_1 \in \mathbb{N}$ depending on T and λ such that $\lambda^{n_1} C_f \geq 2 \max_{e \in E(T)} \mathrm{len}(e)$ so $f^{n_1}(\theta)$ contains an edge.

There exists $n_2 \in \mathbb{N}$ depending on l such that $\lambda^{n_2} \min_{e \in E(T)} \mathrm{len}(e) \geq l$. For $n \geq n_2$, for any $e \in E(T)$, $f^n(e)$ contains a leaf segment of Λ_f^+ with length l.

Then for any $n \geq n_1 + n_2$, $[f^n(\alpha)]$ contains a leaf segment of Λ_f^+ with length l. □

Definition 3.2.4. A lamination Λ is *minimal* in T if all leaves of Λ_T have the same leaf segments up to the action of G, i.e. for every leaf segment $\sigma \subset T$ of Λ_T, for every leaf ℓ in T, there exists $g \in G$ such that $g\sigma \subset \ell$.

Definition 3.2.5. Let Λ be a lamination. Let $T \in \mathcal{D}$. A leaf $\ell \in \Lambda_T$ is *quasi-periodic* if for every $C > 0$ there exists $L > 0$ such that for every subsegment $\sigma \subset \ell$ with $\mathrm{len}(\sigma) = C$, for every subsegment $\gamma \subset \ell$ with $\mathrm{len}(\gamma) > L$, there exists $g \in G$ such that $g\sigma \subset \gamma$.

Remark 3.2.6. If $T \to S$ is a quasi-isometry and Λ_T is minimal, then Λ is minimal in S. Similarly, if a leaf in T is quasi-periodic, then the realization of this leaf in another tree $S \in \mathcal{D}$ is also quasi-periodic. A proof is given in Chapter 2, Remark 2.1.17.

The following is proved in [BFH97], although for a slightly different definition of the stable lamination.

Lemma 3.2.7. *Let $f : T \to T$ be an irreducible train track representative for an automorphism $\phi \in \mathrm{Out}(G)$ with Lipschitz constant λ. Then the stable lamination Λ_f^+ is minimal and its leaves in T are quasi-periodic.*

Proof. First let us prove the minimality. There exists $N \in \mathbb{N}$ such that for all $e, e' \in E(T/G)$, $f^N(e)$ contains an edge in the orbit of e'.

Let $\ell \subset T$ be a leaf of Λ_f^+. Let $k \in \mathbb{N}$. Let us prove that there exists a constant L_k such that every segment of ℓ longer than L_k contains a translate of $f^k(e_0)$ for every $e_0 \in E(T/G)$.

Let σ be a subsegment of ℓ with length at least $L_k := 2\lambda^k \max_{e \in E(T/G)} \mathrm{len}(f^N(e))$. By definition of the stable lamination, there exists $e \in E(T)$ such that $\sigma \subset f^{N+n}(e)$ with $n \geq k$. The segment σ is contained in the concatenation of segments $f^k(f^N(e'))$ for edges $e' \subset f^{n-k}(e)$. Each of these segments contains a translate of $f^k(e_0)$ and is shorter than $\mathrm{len}(\sigma)/2$ so one of them is contained in σ, hence in ℓ. This proves that ℓ contains a translate of every leaf segment of Λ_f^+ contained in $f^k(e_0)$ for any $e_0 \in E(T)$.

This also proves the quasi-periodicity: for all $C > 0$ there exists $k \in \mathbb{N}$ such that every segment I longer than C is contained in $f^k(e)$ for some edge e, and every leaf segment of the leaf ℓ longer than L_k contains a copy of $f^k(e)$ and thus a copy of I. □

Lemma 3.2.8. *Let Λ, Λ' be two distinct minimal closed G-invariant laminations with quasi-periodic leaves. Then for any G-tree T, there exists a bound C_T on the length of leaf segments which are common to both laminations.*

Proof. Let $T \in \mathcal{D}$. By contraposition we prove that if the bound C_T does not exist then Λ and Λ' have the same sets of leaf segments in T. Since they are closed this implies $\Lambda = \Lambda'$.

Suppose that there exist leaf segments $(\eta_n)_{n\in\mathbb{N}}$ with $\eta_n \subset \Lambda_T \cap \Lambda'_T$ for all $n \in \mathbb{N}$, and $\mathrm{len}(\eta_n) \to \infty$. Then for any leaf segment $\sigma \subset \Lambda_T$, there exists $n \in \mathbb{N}$ and $g \in G$ such that $g\sigma \subset \eta_n \subset \Lambda'_T$ so σ is also a leaf segment of Λ'_T. By symmetry Λ_T and Λ'_T have the same leaf segments, so the laminations Λ and Λ' are equal. □

We need the following:

Lemma 3.2.9. *Let $h : T \to T'$ be a (K, C)-quasi-isometry where $K \geq 1, C \geq 0$.*

For every $l > 0$ there exists $L > 0$ depending on l, K, C such that if η is a bi-infinite geodesic, if σ is a subsegment of η with length at least L, then $[h(\sigma)]$ contains a subsegment of $[h(\eta)]$ with length at least ℓ.

Proof. Let $l > 0$. Let η be a bi-infinite geodesic. Let σ be a subsegment of η. Then the length of $[h(\sigma)]$ is at least $K^{-1}\,\mathrm{len}(\sigma) - C$.

The image $h(\eta)$ lies in a $\mathrm{BBT}(h)$-neighbourhood of $[h(\eta)]$, where $\mathrm{BBT}(h) \leq K^2C + C$. The endpoints of $[h(\sigma)]$ are in this neighbourhood so the length of $[h(\sigma)] \cap [h(\eta)]$ is at least $K^{-1}\,\mathrm{len}(\sigma) - C - 2\,\mathrm{BBT}(h)$. Thus by taking $L \geq K(l + C + 2\,\mathrm{BBT}(h))$, if $\mathrm{len}(\sigma) \geq L$, then $[h(\sigma)]$ contains an l-segment of $[h(\eta)]$. □

For a lamination Λ, $T \in \mathcal{D}$ and $C > 0$, a *C-piece* of Λ_T is a leaf segment of Λ_T with length C.

Lemma 3.2.10. *Let $f : T \to T$ and $f' : T' \to T'$ be two train track representatives for a fully irreducible automorphism $\phi \in \mathrm{Out}(G)$. Then the stable laminations Λ^+_f and $\Lambda^+_{f'}$ are equal.*

Proof. We will prove that every leaf segment of $(\Lambda^+_f)_{T'}$ is also a leaf segment of $(\Lambda^+_{f'})_{T'}$: by symmetry we will get the result.

Let $C > 0$. By quasi-periodicity of the leaves of Λ^+_f, there exists $L > 0$ such that every leaf segment of $(\Lambda^+_f)_{T'}$ longer than L contains every orbit of leaf segment of $(\Lambda^+_f)_{T'}$ with length at most C.

Let $h : T \to T'$ be a G-equivariant quasi-isometry.

By Lemma 3.2.9 there exists $L_0 > 0$ such that for every bi-infinite geodesic $\eta \subset T$, for every segment $\sigma \subset \eta$ with length at least L_0, the segment $[h(\sigma)] \cap [h(\eta)]$ has length at least $2L + 2\,\mathrm{BBT}(h)$. Without loss of generality, we may assume $L_0 \geq C_f$.

There exists $g \in G$ be such that $\mathrm{Axis}_{T'}(g)$ is legal for f'. The conjugacy class of g is not pseudo-periodic since $\|\phi^n(g)\|_{T'} \to \infty$ when $n \to \infty$.

The axis of g in T does not have to be legal, however the number of orbits of f-illegal turns under the action of $\langle \phi^n(g) \rangle$ cannot increase when $n \to \infty$: f sends f-legal subsegments to f-legal subsegments and $\mathrm{Axis}_T(\phi^n(g)) = [f^n(\mathrm{Axis}_T(g))]$. Since $\{\|\phi^n(g)\|_T, n \in \mathbb{N}\}$ is unbounded, this implies that there exists $n_0 \in \mathbb{N}$ such that $\mathrm{Axis}_T(\phi^{n_0}(g))$ contains an f-legal subsegment with length L_0. Since $\mathrm{Axis}_{T'}(\phi^{n_0}(g))$ is also f'-legal, up to replacing $g \in \phi^{n_0}(g)$ we can assume $n_0 = 0$.

As in Lemma 3.2.3 there exists $N > 0$ such that for every $e \in E(T')$, $f'^N(e)$ is a leaf segment of $(\Lambda^+_{f'})_{T'}$ with length at least $2L$. The axis of $\phi^N(g)$ in T still contains a legal subsegment with length at least L_0 since $L_0 \geq C_f$. Once again, up to replacing g by $\phi^N(g)$, we may assume that the axis of g in T' can be cut into pieces of $(\Lambda^+_{f'})_{T'}$ with length at least $2L$.

The map h maps $\mathrm{Axis}_T(g)$ to a $\mathrm{BBT}(h)$-neighbourhood of $\mathrm{Axis}_{T'}(g)$. By Lemma 3.2.9 and by definition of L_0, $h(\mathrm{Axis}_T(g)) \subset T'$ contains a leaf segment of $[h(\Lambda^+_f)_T)] = (\Lambda^+_f)_{T'}$ with length at least $2L + 2\,\mathrm{BBT}(h)$ and a subsegment with length at least $2L$ is contained in $\mathrm{Axis}_{T'}(g)$.

Then there exists a segment $\gamma' \subset T'$ of length greater than L that is both a leaf segment $(\Lambda^+_f)_{T'}$ and $(\Lambda^+_{f'})_T$. Thus every leaf segment of $(\Lambda_f)_{T'}$ of length C is a leaf segment of $(\Lambda_{f'})_{T'}$. □

From now on, we can simply refer to the stable lamination as Λ^+_ϕ or simply Λ^+ when the automorphism is obvious. The notation Λ^-_ϕ denotes the unstable lamination $\Lambda^+_{\phi^{-1}}$.

Lemma 3.2.11. *For a fully irreducible automorphism $\phi \in \mathrm{Out}(G)$, the stable lamination and unstable lamination are distinct.*

Proof. Assume by contradiction that $\Lambda^-_\phi = \Lambda^+_\phi$. See Figure 3.5.

Let $f : T \to T$ be a train track representative for ϕ and $f_- : T_- \to T_-$ be a train track representative for ϕ^{-1}. Let $\tau : T \to T_-$ be a G-equivariant quasi-isometry.

There exists $g \in G$ loxodromic whose axis in T is f-legal. By Lemma 3.2.9 there exists $L > 0$ such that for every leaf segment σ of $(\Lambda^+)_T$ longer than L, $[\tau(\sigma)]$ contains a leaf segment of $(\Lambda^+)_{T_-}$ longer than $2C_{f_-}$.

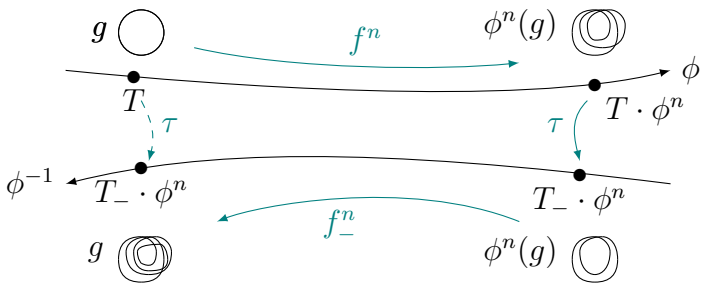

Figure 3.5 – When applying f^n and then f_-^n, the translation length of g increases if n is big enough

Let $n_0 \in \mathbb{N}$ be such that for any edge $e \in T$, $\mathrm{len}(f^n(e)) \geq 2C_f$. Let $n \geq n_0$. Then every $e \in \mathrm{Axis}_T(\phi^n(g))$ contributes a leaf segment of $(\Lambda^+)_{T_-} = (\Lambda^-)_{T_-}$ longer than $2C_f$ in $\mathrm{Axis}_{T_-}(\phi^n(g))$. By Lemma 3.1.13, the images $\beta_e := [\tau(f^n(e))]$ for different edges e contain subsegments β'_e obtained from β_e by cutting out the $\frac{C_{f_-}}{2}$-neighbourhood of the endpoints. The subsegments β'_e satisfy the following: for any edge e, for any $m \in \mathbb{N}$, $f_-^m(\beta'_e) \subset \mathrm{Axis}_{T_-}(\phi^{n-m}(g))$ and for any other edge e', $f_-^m(\beta'_e) \cap f_-^m(\beta'_{e'}) = \varnothing$.

This implies that $[f_-^n(\mathrm{Axis}_{T_-}(\phi^n(g)))] = \mathrm{Axis}_{T_-}(g)$ contains two disjoint leaf segments $f_-^n(\beta'_e)$, $gf_-^n(\beta'_e) = f_-^n(\beta'_{\phi^n(g)e})$ longer than $\lambda^n C_f$.

Then we must have $\|g\|_{T_-} \geq \lambda^n C_f$, which is a contradiction for n sufficiently big. □

3.3 Laminations and simple elements of G

3.3.1 Simple elements, simple pairs and Whitehead graphs

Definition 3.3.1. A loxodromic element $g \in G$ is *simple* if it is contained in a proper cyclic factor of G.

A pair of elements $g, h \in G$ is *simple* if there exists cyclic factors H_g, H_h such that $g \in H_g, h \in H_h$ and a graph Γ of cyclic groups with $\pi_1(\Gamma) \simeq G$, with disjoint subgraphs Γ_g, Γ_h such that $\pi_1(\Gamma_g) \simeq H_g, \pi_1(\Gamma_h) \simeq H_h$.

In that case H_g, H_h belong to a proper *system of cyclic factors*, i.e. collection of conjugacy class of cyclic factors which can be simultaneously seen in some graph of groups as the fundamental groups of disjoint subgraphs.

Lemma 3.3.2. *Suppose $b_1(G) \geq 3$. For any tree $T \in \mathcal{D}$, the candidates in T are simple.*

Proof. If $b_1(G) \geq 3$ then every candidate g for T avoids at least an orbit of edges $G \cdot e$. Then $T \setminus G \cdot e$ is a proper subforest of G which contains the axis of g. It defines a cyclic factor containing g, so g is simple. □

Lemma 3.3.3. *Suppose $b_1(G) \geq 3$. Suppose $g, h \in G$ are candidates in $T \in \mathcal{D}$. There exists a candidate $k \in G$ such that $\{k, g\}$ and $\{k, h\}$ are both simple.*

Proof. Since $b_1(\Gamma) \geq 3$ and g, h are candidates, neither of their axes crosses every orbit of edges in T.

Let e_g, e_h be edges in Γ such that $\pi(\mathrm{Axis}_T(g))$ avoids e_g and $\pi(\mathrm{Axis}_T(h))$ avoids e_h, where $\pi : T \to \Gamma$ is the quotient map. Since $b_1(g) \geq 3$ the graph $\Gamma' := \Gamma \setminus \{e_g, e_h\}$, which may be disconnected, has a connected component with first Betti number $b_1(\Gamma') \geq 1$. There exists an element $k \in G$ whose axis in T is in a lift of Γ'. Then $\mathrm{Axis}_T(k)$ crosses neither e_g nor e_h so:

— the axes of g, k are in $\Gamma \setminus e_g$, so $\{k, g\}$ is simple

— the axes of h, k are in $\Gamma \setminus e_h$, so $\{k, h\}$ is simple

Thus k is the element of G that we were looking for. □

Definition 3.3.4. Let ℓ be a bi-infinite geodesic in $T \in \mathcal{D}$. A *turn* in ℓ is a pair $\{e, e'\} \subset E(T)$ of distinct edges such that $o(e) = o(e')$ and $e \cup e' \subset \ell$.

Definition 3.3.5. Let $\mathcal{G}$ be a collection of bi-infinite geodesics of some $T \in \mathcal{D}$. Let $v \in V(T)$. The Whitehead graph $W := \mathrm{Wh}_T(\mathcal{G}, v)$ is the following graph:

— vertices of W are edges of T with origin v

— there is an edge $e - e'$ in W if there exists $\ell \in \mathcal{G}$ and $g \in G$ such that $g \cdot \ell$ contains both e and e', i.e. if $\{e, e'\}$ is a turn crossed by ℓ

Remark 3.3.6. For $\mathcal{G}, T, v$ as in the definition, we have $\mathrm{Wh}_T(\mathcal{G}, v) = \mathrm{Wh}_T(G \cdot \mathcal{G}, v)$.

Examples 3.3.7. The two main examples, which we will both use in this paper, are the following.

1. $\mathcal{G} = \{\mathrm{Axis}_T(hgh^{-1}), h \in G\}$ is the collection of axes of all conjugates of some $g \in G$. In that case we write $\mathrm{Wh}_T(\mathcal{G}, v) = \mathrm{Wh}_T(g, v)$.
2. let $f : T \to T$ be a train track representative for a fully irreducible automorphism and let $S \in \mathcal{D}$, let $\mathcal{G} = (\Lambda^+)_S$. Since all leaves in Λ^+ have the same subsegments in S, for any leaf $\ell \in (\Lambda^+)_S$, for any $v \in V(S)$, $\mathrm{Wh}_S(\Lambda^+, v) = \mathrm{Wh}_S(\ell, v)$.

The interest of Whitehead graphs is that they help understanding cyclic factors. In Chapter 1 we prove the following theorem (Theorem 1.2.14):

Theorem 3.3.8. — *Let $g \in G$ be a loxodromic element. Then g is simple if and only if for every $T \in \mathcal{D}$ there exists $v \in T$ such that $\mathrm{Wh}_T(g, v)$ is disconnected or has a cut vertex, i.e. a vertex $p \in \mathrm{Wh}_T(g, v)$ such that $\mathrm{Wh}_T(g, v) \setminus \{p\}$ is disconnected.*

— *Let $g, h \in G$ be loxodromic elements. Then $\{g, h\}$ is simple if and only if for every $T \in \mathcal{D}$ there exists $v \in T$ such that $\mathrm{Wh}_T(\{g, h\}, v)$ is disconnected or has a cut vertex.*

The stable lamination of an automorphism $\phi \in \mathrm{Out}(G)$ is *carried* by a cyclic factor H if and only if there exists $T \in \mathcal{D}$ such that for every leaf $\ell \in (\Lambda^+)_T$ there exists a translate of the minimal subtree T_H which contains ℓ.

Minimality of these laminations imply that the stable lamination is carried by H if and only if there exists a leaf $\ell \in (\Lambda^+)_T$ such that ℓ is contained in T_H. For all $S \in T$, there exists a quasi-isometry $T \to S$ and it implies that the realization of ℓ in S is in the subtree S_H, so the fact that Λ^+ is carried by H can be seen in every $S \in \mathcal{D}$. These facts are proved in Chapter 2.

Lemma 2.2.5 of Chapter 2 implies:

Proposition 3.3.9. *Suppose $\phi \in \mathrm{Out}(G)$ is a fully irreducible automorphism with irreducible train track representative $f : T \to T$. Then no leaf of the stable lamination $(\Lambda^+_\phi)_T$ is carried by a cyclic factor.*

3.3.2 Long segments of laminations in axes of elements of G

In this section, we assume that $\phi \in \mathrm{Aut}(G)$ is a fully irreducible automorphism such that both ϕ and ϕ^{-1} have train track representatives.

The following lemma is a transposition of Lemma 1.2.17 from Chapter 1. The original lemma gives a link between the Whitehead graph of the axis of a loxodromic element $g \in G$ in a tree S and the existence of a tree $\hat{S}$ where $\hat{S} \to S$ is either a fold or a collapse which induces an isometry $\mathrm{Axis}_{\hat{S}}(g) \to \mathrm{Axis}_S(g)$. The proof of the lemma does not use the specific fact that $\mathrm{Axis}_S(g)$ is an axis of an element and could actually work with any bi-infinite geodesic. In particular, it can be transposed to laminations:

Lemma 3.3.10. *Let $S \in \mathcal{D}$ such that no edge in S/G is a loop. Let ℓ^+ be a leaf of the stable lamination Λ^+ and ℓ^- be a leaf of the unstable lamination Λ^-. The following are equivalent:*

— *There exists a vertex* $v \in V(S)$ *such that* $\mathrm{Wh}_S(\{\ell^+, \ell^-\}, v)$ *is disconnected or has a cut vertex*

— *There exists a tree* $\hat{S} \in \mathcal{D}$ *and a non-injective map* $\pi := \hat{S} \to S$ *such that if* $\hat{\ell}^+, \hat{\ell}^-$ *are the leaves in* $\hat{S}$ *corresponding to* ℓ^+, ℓ^- *then* π *induces isometries* $\hat{\ell}^+ \to \ell^+$ *and* $\hat{\ell}^- \to \ell^-$.

Remarks 3.3.11. 1. The assumption that S has no loop is not especially restrictive: in fact, up to subdividing all loops before applying the lemma, we may assume that S has no loop. The tree $\hat{S}$ produced by the lemma does not have any loop either.

2. A non-injective G map $\hat{S} \to S$ sending vertex to vertex and edge to edge is a composition of collapses and folds (see [BF91]). In particular we can assume that the map π given by the lemma is either a collapse or a fold.

Proposition 3.3.12. *There exists* $S \in \mathcal{D}$ *such that for every* $v \in V(S)$*, the Whitehead graph* $\mathrm{Wh}_S(\Lambda^+ \cup \Lambda^-, v)$ *is connected without cut vertex.*

Proof. Let T be the initial train track representative for ϕ. We will change T gradually using Lemma 3.3.10.

By Proposition 3.3.9 from Chapter 2, the Whitehead graph $\mathrm{Wh}_T(\Lambda^+, v)$ is connected for every $v \in S$. The graph $\mathrm{Wh}_T(\Lambda^+ \cup \Lambda^-, v)$ has even more edges so it is also connected.

If an edge of T is a loop then subdivide it so that Lemma 3.3.10 applies. Endow T with the combinatorial metric, i.e. give each edge the length 1. This does not change the Whitehead graphs.

Suppose there exists a Whitehead graph in T which has a cut point. By applying Lemma 3.3.10 and Remark 3.3.11, we can construct a sequence

$$\dots T_n \to T_{n-1} \to \cdots \to T_0 = T$$

where each map $T_i \to T_{i-1}$ is either a collapse or a fold whose restriction to the leaves of both laminations are isometric. The construction of the sequence stops when we find n such that every Whitehead in T_n has no cut vertex.

The maps $T_i \to T_{i-1}$ are in fact not collapses, since a collapsed edge in T_i would not be crossed by any leaf of the lamination, contradicting Proposition 3.3.9. They are folds.

We want to prove that the sequence above cannot be infinite. By contradiction, assume it is infinite. In Chapter 1, Lemma 1.1.25 we proved that the number of orbits of edges

of the trees of the sequence built by iterating Lemma 3.3.10 has to go to infinity in that case.

The first Betti number $b_1(T_n)$ is constant. Recall that for a connected graph Γ with V vertices and A edges we have $b_1(\Gamma) = A - V + 1$.

Recall that a *big* vertex stabilizer is a vertex stabilizer which does not fix any edge in some (equivalently any) reduced tree. By Lemma 3.1.1 there is a bound on the number of vertices of valence 1 in T_n/G. In fact, the associated vertex groups are big since trees in $\mathcal{D}$ are minimal. Since there exist finitely many conjugacy classes of big vertex stabilizers, this gives a bound on the number of vertices of valence 1 in T_n/G.

Let A_n, V_n be the number of edges and vertices in T_n/G. For every $v \in T_n$ denote by $\mathrm{val}(v)$ the valence of v. Then we have

$$2V_n + 2b_1(\Gamma) - 2 = 2A_n = \sum_{v \in T_n/G} \mathrm{val}(v)$$

Therefore

$$2b_1(\Gamma) - 2 = \sum_{v \in T_n/G} (\mathrm{val}(v) - 2)$$

The only negative terms in the sum correspond to valence 1 vertices so there is a lower bound on their sum. This implies that there is a bound on the number of vertices with valence ≥ 3.

As a result, since the number of edges in T_n when n goes to infinity is unbounded, the number of vertices of valence 2 in T_n/G is unbounded.

Let $v \in T_n/G$ be a vertex of valence 2. Let l_1, l_2 be the labels at v. If $|l_1| > 1$ and $|l_2| > 1$ then the stabilizers of vertices in the orbit of v are big. Thus the number of vertices of valence 2 with both labels distinct from ± 1 is bounded by $m(G)$. For every other vertex of valence 2, one of the labels is 1 or -1.

A *topological edge* in T_n/G is a connected component of

$$\Gamma \setminus \{v \in V(\Gamma) / \mathrm{val}(v) \neq 2 \text{ or no label at } v \text{ is } \pm 1\}$$

There is a bound B, independent of n, on the number of topological edges in T_n/G. Since $\mathrm{vol}(T_n/G)$ is unbounded there is no bound on the length of topological edges.

A subsegment $\sigma := e_0, \ldots, e_k$ of a topological edge is *increasing* if for any $i \in \{1, \ldots, k\}$, $G_{o(e_i)} = G_{e_i}$. It is *decreasing* if for any $i \in \{0, \ldots, k-1\}$, $G_{t(e_i)} = G_{e_i}$. A

subsegment with a single edge is both increasing and decreasing. This can be understood efficiently with labels: σ is increasing (resp. decreasing) if $\lambda(e_i) = \pm 1$ (resp. $\lambda(\bar{e}_i) = \pm 1$) for all $i \in \{1, \dots, k\}$.

We will now prove that any topological edge of T_n/G can be cut into at most $2m(G)+1$ subsegments which are either increasing or decreasing. See Figure 3.6 for an example.

Let $\sigma := e_0, \dots, e_k$ be a topological edge. Cut it into subsegments by the following process. Let $\sigma_1 := e_0, \dots, e_{i_1}$ be the maximal decreasing prefix of σ: it has at least one edge. Let σ_2 be the maximal increasing prefix of $\sigma \setminus \sigma_1$. The label $\lambda(\bar{e}_{i_1})$ must be nonzero unless $\sigma_1 = \sigma$, so since σ is a topological edge, $\lambda(e_{i_1+1}) = \pm 1$ so σ_2 also has at least one edge. Continue this procedure to construct an alternating sequence of disjoint decreasing and increasing subsegments.

Write σ as the concatenation $\sigma_1, \dots, \sigma_N$ of subsegments. We claim that whenever an increasing subsegment is followed by a decreasing subsegment, the last edge of the former has a big stabilizer, and no edge in the latter does. Thus the number of increasing subsegments in σ is bounded by $m(G) + 1$.

Suppose σ_j is an increasing subsegment followed by a decreasing subsegment σ_{j+1}. Let e_i be the last edge of σ_j. Then by maximality of σ_j, $\lambda(e_{i+1}) \neq \pm 1$. Besides, σ_1 is a decreasing subsegment so there is a decreasing subsegment before σ_j. Let $e_{i'}$ be its last edge. We have $i' < i$ and $\lambda(\bar{e}_{i'}) \neq \pm 1$. For all $p \in \{i'+1, i\}$, $\lambda(e_p) = \pm 1$ so σ_j is collapsible and collapses to a vertex v with labels $\lambda(e_{i+1})$ on the right, $\prod_{i' \leq p \leq i} \lambda(e_p)$ on the left. Both labels are not ± 1 so the vertex group associated to v is big. It is also the edge group associated to e_i.

For a topological edge of length k, at least one the maximal topological edges is longer than $\frac{k}{2m(G)+2}$. Thus there is no bound on the maximal length of half topological edges when n increases.

Suppose σ is an increasing subsegment of a topological edge with length k in T_n/G. Let us prove that there exists a leaf of Λ^+ and a leaf of Λ^- which overlap along a segment with length k. Let $\pi_n : T_n \to T_n/G$ be the quotient map.

Write σ as the concatenation $e_0, \dots, e_k$ with $\lambda(e_i) = \pm 1$ for all $i \in \{1, \dots, k\}$ (see Figure 3.7). Let $w = t(e_k), v = o(e_0)$.

The subsegment σ lifts in T to a subforest Y_σ. Let $\tilde{\sigma}$ be a connected component of $Y_\sigma \backslash \pi_n^{-1}(v)$. It is a finite rooted tree with root $\tilde{w}$ which is a lift of w. For every $i \in \{1, \dots, k\}$, any lift of $\tilde{e}_i$ is an edge of T which points towards $\tilde{w}$. The terminal vertices of $\tilde{\sigma}$, other than possibly $\tilde{w}$, are the lifts of v.

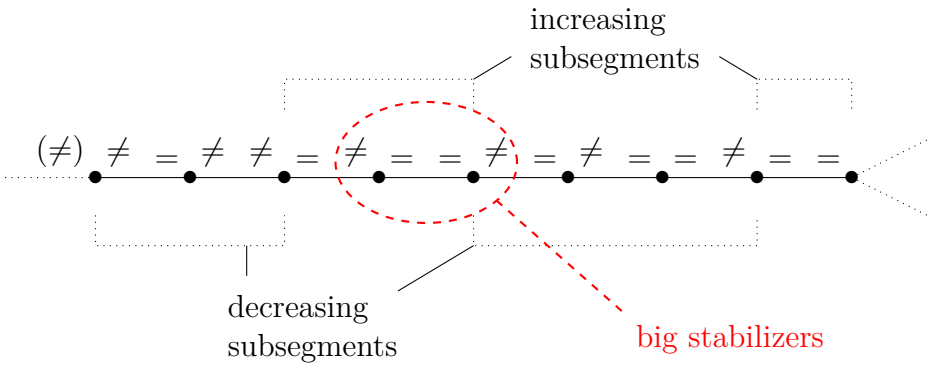

Figure 3.6 – General form of a topological edge and monotonous subsegments

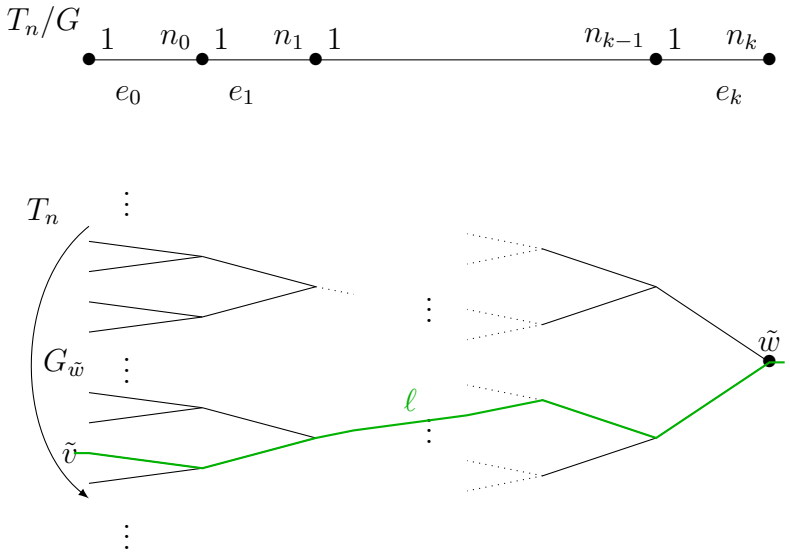

Figure 3.7 – Pre-image of an increasing subsegment

The vertex group $G_{\tilde{w}}$ acts transitively on the set of lifts of v.

Let ℓ be a leaf of Λ^+. There is a translate of ℓ which crosses a lift of e_k. Thus it contains $[\tilde{v}, \tilde{w}]$ where $\tilde{v}$ is a lift of v. By transitivity of the action of $G_{\tilde{w}}$, for every $\tilde{v}' \in \pi_n^{-1}(v)$, the segment $[\tilde{v}', \tilde{w}]$ is contained in a leaf of Λ^+.

Similarly, if ℓ_- is a leaf of Λ^-, there exists a translate of ℓ_- crossing $\tilde{e}_k$ and by translating further by an element of $G_{\tilde{w}}$ we can make sure that it crosses $[\tilde{v}', \tilde{w}]$ for any arbitrary $\tilde{v}' \in \pi_n^{-1}(v)$. Then the translates of ℓ and ℓ_- overlap on a length at least k.

The same proof can be transposed to the case of a decreasing subsegment.

Since there is no bound on the length of monotonous subsegments, for every $L > 0$

there exists $n \in \mathbb{N}$ such that T_n contains an increasing or decreasing subsegment σ longer than L. Let $\tilde{\sigma}$ be a lift for σ in T_n. There exist leaves ℓ^+, ℓ^- of the stable and unstable laminations which both cross $\tilde{\sigma}$. Therefore the leaves ℓ^+ and ℓ^- overlap on a length bigger than L.

The maps $T_n \to T$ for $n \in \mathbb{N}$ are isometric in restriction to the leaves of the laminations. Thus there is no bound on the length of common subsegments of both laminations in T, so by Lemma 3.2.8 the laminations are equal. This is a contradiction to Lemma 3.2.11. □

Proposition 3.3.13. *Let $\phi \in \mathrm{Out}(G)$ be a train track automorphism. Let $f : T \to T$ be a train track representative for ϕ. Let $\mathcal{L}_f$ be an axis for ϕ in $\mathcal{D}$ passing through T.*

There exists $L > 0$ such that for any $T_t \in \mathcal{L}_f$:

(i) If g is a simple loxodromic element in G, then $\mathrm{Axis}_{T_t}(g)$ cannot simultaneously contain an L-piece of Λ^+ and an L-piece of Λ^-.

(ii) If g, h are simple loxodromic elements such that $\mathrm{Axis}_{T_t}(g)$ contains an L-piece of Λ^+ and $\mathrm{Axis}_{T_t}(h)$ contains an L-piece of Λ^-, then the pair $\{g, h\}$ is not simple.

Proof. Let S be a tree obtained with Proposition 3.3.12, i.e. such that for every $v \in V(S)$ the Whitehead graph $\mathrm{Wh}_S(\Lambda^+ \cup \Lambda^-, v)$ is connected without cut vertex. Note that for any $n \in \mathbb{Z}$, $S \cdot \phi^n$ has the same property since Λ^+, Λ^- are ϕ-invariant.

By quasi-periodicity of leaves, there exists a constant $L_1 > 0$ such that any leaf segment of Λ^+ in S (resp. Λ^-) longer than L_1 contains all turns in Λ^+ (resp. Λ^-). Suppose $\mathrm{Axis}_S(g)$ contains an L_1-piece of Λ^+ and an L_1-piece of Λ^-, then $\mathrm{Wh}_S(g, v)$ is connected without cut vertex for all $v \in V(S)$. By theorem 3.3.8 g is not simple. This proves assertion (i) in the specific case where the tree is S: now we would like to prove it for T_t in the axis of ϕ.

By Remark 3.1.28 there exist constants $K > 1, C > 0$ such that for every $t \in [0, \log(\lambda)]$ there exists a G-equivariant (K, C)-quasi-isometry $h_t : T_t \to S$. There is an upper bound $B \geq 0$, depending only on (K, C), on $\mathrm{BBT}(h_t)$.

Let L_0 be the constant of Lemma 3.2.9 for the quasi-isometry constants K, C and $\ell = L_1 + 2B$. If a fundamental domain γ of g in T_t contains a leaf segment σ with $\mathrm{len}(\sigma) \geq L_0$ then $[h_t(\gamma)]$ contains a leaf segment with length at least $L_1 + 2B$, and a fundamental domain of g in S contains an L_1-piece of the lamination. This works for both laminations.

Let $L > \max\{L_0, 2KL_1\}$. Suppose $\mathrm{Axis}_{T_t}(g)$ contains leaf segments σ_+, σ_- of the stable and unstable laminations, both longer than L. Up to replacing g with g^l for some $l \in \mathbb{N}$ we can suppose a fundamental domain for g contains both σ_+ and σ_-. Lemma 3.2.9 ensures

that $\mathrm{Axis}_S(g)$ contains L_1-pieces of both laminations, therefore implying that g^l hence g is not simple.

Finally suppose $t \notin [0, \log(\lambda)]$. There exists $n \in \mathbb{Z}$ such that $T_{t'} := T_t \cdot \phi^n$ with $t' = t + n\log(\lambda) \in [0, \log(\lambda)]$. Then $h_{t'} : T_{t'} \to S$ induces a G-equivariant (K, C)-quasi-isometry $T_t \to S \cdot \phi^{-n}$. With the same arguments as above we come to the same result, with the same constant L. This proves (i).

The proof of (ii) is analogous. We just proved that there exists $L > 0$, such that for $T_t \in \mathcal{L}_f$ there exists n such that for any $g \in G$, if $\mathrm{Axis}_{T_t}(g)$ contains an L-piece of any lamination then $\mathrm{Axis}_{S\cdot\phi^n}(g)$ contains an L_0-piece of the same lamination. Applying this to g with the stable lamination and h with the unstable lamination, we get that $\mathrm{Wh}_{S\cdot\phi^n}(\{g, h\}, v)$ is connected without cut vertex for any $v \in V(S)$ and therefore $\{g, h\}$ is not simple. □

3.4 Legality

Let G be a GBS group with first Betti number $b_1(G) \geq 3$.

In this section we fix a pseudo-atoroidal fully irreducible automorphism $\phi \in \mathrm{Out}(G)$ with a train track representative $f : T \to T$ and a train track representative $f_- : T_- \to T_-$ for ϕ^{-1}. The goal is to study the evolution of $\|\phi^n(g)\|_T$, $\|\phi^n(g)\|_{T_-}$ when $n \to \pm\infty$ for $g \in G$.

The following three lemmas prove an analogue of Lemmas 2.9 and 2.10 in [BFH97]. The point of view differs a little since we state the results in the trees and not in the quotient graphs. A notable difference which is caused by non-trivial edge stabilizers is the fact that a concatenation of pINPs is not always a Nielsen path, since it might only be pre-periodic. The statements also differ a little for technical reasons.

Lemma 3.4.1. *Let $\phi \in \mathrm{Aut}(G)$ be an automorphism with a train track representative $f : T \to T$. For every $C > 0$ there exists $M \in \mathbb{N}$ such that for any edge path $\sigma \subset T$, one of the following holds:*

(a) $[f^M(\sigma)]$ is legal

(b) $[f^M(\sigma)]$ contains a legal segment of length $\geq C$ between two illegal turns

(c) $[f^M(\sigma)]$ has fewer illegal turns than σ

(d) *σ is a concatenation $\eta_0 \cdot \eta_1 \cdot \cdots \cdot \eta_{k+1}$ for some $k \geq 1$ where η_0, η_{k+1} are legal subpaths, and for $1 \leq i \leq k$ the path $f^M(\eta_i)$ is a periodic indivisible Nielsen path, and turns at the concatenation points are legal.*

Proof. Without loss of generality we may assume $C > C_f$. Then if $\sigma := [y, y']$ contains a legal path with length C then for any $n \in \mathbb{N}$, $[f^n(\sigma)]$ also does, by definition of the critical constant.

Let $M \in \mathbb{N}$ be a big enough integer, to be determined later. Suppose there exists a path $\sigma \subset T$ such that both (a), (b) and (c) fail. Since $[f^M(\sigma)]$ cannot have more illegal turns than σ, it has exactly the same number of illegal turns $k \geq 1$. There exist maximal legal subsegments $\gamma_0, \ldots, \gamma_k$ such that σ is the concatenation $\gamma_0 \cdot \cdots \cdot \gamma_k$. Since (a) fails we have $\mathrm{len}(\gamma_i) \leq C$ for every $i \in \{1, \ldots, k-1\}$.

The map f maps legal segments to legal segments and since the number of illegal turns is constant, for every $n \leq M$, there is a unique decomposition $[f^n(\sigma)] = \gamma_0^n \cdot \gamma_1^n \cdot \cdots \cdot \gamma_k^n$ into maximal legal subsegments. We have $\mathrm{len}(\gamma_i^n) \leq C_f$ for all $i \in \{1, \ldots, k-1\}$.

There are finitely many orbits of edge paths of T with length at most $2C_f$. Let N be the number of orbits of such subpaths.

Let $i \in \{2, \ldots, k-1\}$. There exists $p_i \leq N$ and $g_i \in G$ such that $\gamma_{i-1}^N \cdot \gamma_i^N = g_i \left(\gamma_{i-1}^{N+p_i} \cdot \gamma_i^{N+p_i}\right)$.

There also exists $p_1 \leq N$ and $g_1 \in G$ such that the restrictions of $\gamma_0^N \cdot \gamma_1^N$ and $g_1 \left(\gamma_0^{N+p_1} \cdot \gamma_1^{N+p_1}\right)$ to a $2C_f$-neighbourhood of the illegal turn are equal. Similarly define $p_k \in \mathbb{N}, g_k \in G$ such that restrictions of $\gamma_{k-1}^N \cdot \gamma_k^N$ and $g_k \left(\gamma_{k-1}^{N+p_k} \cdot \gamma_k^{N+p_k}\right)$ coincide.

By taking the smallest common multiple of all p_i for $i \in \{1, \ldots, k\}$ we may replace p_i by some $P \in \mathbb{N}$ which does not depend on i.

Thus for every $i \in \{2, \ldots, k-1\}$, we have $\gamma_{i-1}^N \cdot \gamma_i^N \subset g_i[f^P(\gamma_{i-1}^N \cdot \gamma_i^N)]$. The same holds for restrictions to a $2C_f$-neighbourhood of the illegal turn for $i \in \{1, k\}$. Note that this implies that there exists no $n \in \mathbb{N}$ such that γ_0^n or γ_k^n vanish.

For any $i \in \{1, \ldots, k\}$ the path γ_i^N contains a unique point x_i such that $x_i = g_i f^P(x_i)$, and γ_0^N contains a unique point x_0 such that $x_0 = g_1 f^P(x_0)$. The point x_0 might be equal to y, x_k might be equal to y'. Note that $g_i \gamma_i^{N+P} = g_{i+1} \gamma_i^{N+P} = \gamma_i^N$ so for $i \in \{0, \ldots, k-1\}$ actually $g_{i+1} f^P(x_i) = g_i f^P(x_i) = x_i$. Then for any $i \in \{1, \ldots, k\}$ the path $[x_{i-1}, x_i]$ is a periodic indivisible Nielsen path.

For any $i \in \{0, \ldots, k\}$ there exists a unique point $y_i \in \sigma$ such that $f^N(y_i) = x_i$. The points $y_0, \ldots, y_k$ subdivide σ into $k+2$ subsegments. Define $\eta_i = [y_i, y_{i+1}]$ for $i \notin \{0, k+1\}$, $\eta_0 = [y, y_0]$ and $\eta_{k+1} = [y_k, y']$: we just proved that η_i is a pre-Nielsen path for $1 \leq i \leq k$.

The other subpaths η_0 and η_{k+1} are legal, and the other subpaths are pre-Nielsen paths, so σ satisfies (d).

The integers N and P only depend on T. If $M \geq N + P$ we proved that for any path σ such that (a), (b) and (c) fail, (d) holds. □

The following result, which is a key for Lemma 3.4.3, implies that when neither Case (a), Case (b) nor Case (d) of Lemma 3.4.1 occur, the decrease of the number of illegal turns is a definite proportion of the length of the segment.

Lemma 3.4.2. *Let $\phi \in \mathrm{Aut}(G)$ with a train track representative $f : T \to T$. Let $C > C_f$. Let $M_0 \geq 1$ be the corresponding integer given by Lemma 3.4.1. There exists $p \in \mathbb{N}$ with the following property. Let $M := pM_0$. There exists $K < 1$ and $K' \geq 0$ such that for any loxodromic $g \in G$, for any $\sigma \subset \mathrm{Axis}_T(g)$, there exists segments α, β, σ' such that $[f^M(\sigma)] = \alpha \cdot \sigma' \cdot \beta$ with*

- *$\mathrm{len}(\alpha) \leq K'/2, \mathrm{len}(\beta) \leq K'/2$*
- *$\sigma' \subset \mathrm{Axis}_T(\phi^M(g))$ or $\mathrm{len}(\sigma') = 0$*
- *either σ' contains a legal subsegment with length greater than C, or*

$$\mathrm{len}(\sigma') \leq K\,\mathrm{len}(\sigma) + K'$$

Moreover, if σ contains a legal subsegment longer than C, then so does σ'.

Proof. Let $C > C_f$. Let $g \in G$ and $\sigma \subset \mathrm{Axis}_T(g)$. Let M_0 be the constant of Lemma 3.4.1 for C.

Let $m \in \mathbb{N}$ be the maximal number of pINPs which may be concatenated in T or T_-, from Lemma 3.1.18.

If $\mathrm{len}([f^{M_0}(\sigma)]) \leq K'_0 := 2\,\mathrm{BBT}(f^{M_0}) + 2mC$, then define $\alpha := [f^{M_0}(\sigma)]$, and σ' and β as single points such that $[f^{M_0}(\sigma)] = \alpha \cdot \sigma' \cdot \beta$. These subsegments satisfy the statement for any choice of p, with $\mathrm{len}(\sigma') = 0$.

Suppose $\mathrm{len}([f^{M_0}(\sigma)]) > K'_0$. Define $\sigma_0 := [f^{M_0}(\sigma)] \cap \mathrm{Axis}_T(g)$, which is not empty, and let α_0, β_0 be the remaining subsegments. Note that if σ contains a legal subsegment θ with length greater than $C > C_f$, then $[f^{iM_0}(\sigma) \cap \mathrm{Axis}_T(\phi^{iM_0}(g)) \cap f^{iM_0}(\theta)$ is longer than C for all $i \geq 1$, hence the last statement.

Let $\mathrm{nbl}(\theta)$ denote the number of maximal legal subsegments in a segment θ. We will now prove that if σ_0 contains no legal subsegment longer than C, then its number of maximal legal subsegments $\mathrm{nbl}(\sigma)$ decreases.

Write $\sigma := \theta_0 \cdot \theta_1 \cdot \cdots \cdot \theta_n$ where each subsegment θ_i except θ_0 has $m+2$ maximal legal subsegments. Let $i \in \{1, \ldots, n\}$. Apply Lemma 3.4.1 to θ_i. Case (d) cannot happen since θ_i has $m+1$ illegal turns. If Case (a) happens then Case (c) also happens: $\mathrm{nbl}([f^{M_0}(\theta_i)]) < \mathrm{nbl}(\theta_i) = m+1$. Suppose Case (b) happens: either $[f^{M_0}(\theta_i)] \cap \mathrm{Axis}_T(\phi^{M_0}(g))$ contains a legal subsegment with length C, or an illegal turn of θ_i is sent outside $\mathrm{Axis}_T(\phi^{M_0}(g))$, in which case $\mathrm{nbl}([f^{M_0}(\theta_i)] \cap \mathrm{Axis}_T(\phi^{M_0}(g))) < \mathrm{nbl}(\theta_i)$.

Thus if σ_0 does not contain any legal subsegment longer than C, then

$$\begin{aligned}\mathrm{nbl}(\sigma_0) &\leq \mathrm{nbl}(\theta_0) + \mathrm{nbl}([f^{M_0}(\theta_1)] \cap \mathrm{Axis}_T(\phi^{M_0}(g))) + \cdots + \mathrm{nbl}([f^{M_0}(\theta_n)] \cap \mathrm{Axis}_T(\phi^{M_0}(g))) \\ &\leq \mathrm{nbl}(\theta_0) + (\mathrm{nbl}(\theta_1) - 1) + \cdots + (\mathrm{nbl}(\theta_n) - 1) \\ &\leq \mathrm{nbl}(\sigma) - n\end{aligned}$$

where $n = \left\lfloor \frac{\mathrm{nbl}(\sigma)}{m+1} \right\rfloor$.

Therefore, with $k := (1 - \frac{1}{m+1})$, we obtain

$$\mathrm{nbl}(\sigma_0) \leq k \, \mathrm{nbl}(\sigma)$$

This can be iterated as long as $[f^{iM_0}(\sigma)] \cap \mathrm{Axis}_T(\phi^{iM_0}(g))$ contains no legal subsegment of length C by applying the same argument to σ_1 instead of σ, creating a decreasing subsequence $\sigma_1, \sigma_2, \ldots$ of σ and increasing sequences $\alpha_1, \alpha_2, \ldots$ and $\beta_1, \beta_2, \ldots$. Then for $p \geq 1$, $[f^{pM_0}(\sigma)]$ can be cut into subsegments $\alpha_p \cdot \sigma_p \cdot \beta_p$ with

— $\mathrm{len}(\alpha_p), \mathrm{len}(\beta_p) \leq K_0' \sum_{i=0}^{p-1} \lambda^i =: K_p'$

— $\sigma_p \subset \mathrm{Axis}_T(\phi^{pM_0}(g))$

— $\mathrm{nbl}(\sigma_p) \leq k^p \, \mathrm{nbl}(\sigma)$

Now convert this result into lengths: we obtain

$$\mathrm{len}(\sigma_p) \leq \frac{Ck^p}{l_{\min}} \mathrm{len}(\sigma)$$

where $l_{\min} := \min_{e \in E(t)} \mathrm{len}(e)$.

Choose p such that $\frac{Ck^p}{l_{\min}} < 1$. Let $K := \frac{Ck^p}{l_{\min}}$ and $K' := K_p'$.

Finally we get

$$\mathrm{len}(\sigma) \leq K \, \mathrm{len}(\sigma) + K'$$

thus we obtain the lemma with $\alpha := \alpha_p$, $\beta := \beta_p$, $\sigma' := \sigma_p$. □

Lemma 3.4.3. *Recall that ϕ is pseudo-atoroidal. Let $h : T \to T_-$ be a Lipschitz G-equivariant map, sending vertex to vertex, and edge to edge path.*

For every $C > 0$ there exists $N \in \mathbb{N}$ and $L > 0$ such that for any $g \in G$, for any geodesic $\sigma \subset \mathrm{Axis}_T(g)$ with length greater than L and possibly infinite, if $\sigma' := [h(\sigma)]$, then one of the followings holds:

(A) $[f^N(\sigma)] \cap \mathrm{Axis}_T(\phi^N(g))$ contains a legal segment of length $> C$

(B) $[f_-^N(\sigma')] \cap \mathrm{Axis}_{T_-}(\phi^{-N}(g))$ contains a legal segment of length $> C$.

Proof. Let M be the constant from Lemma 3.4.2. Assume

$$C \geq \max\left\{C_f, C_{f_-}\right\}$$

We will suppose by contradiction that the lemma fails for $N := Mi$ with i sufficiently big. We will take a segment σ in T for which both (A) and (B) fail, and show that this assumption leads to a contradiction in the following sense. For $j \in \{0, \dots, i\}$ the segment $[f^{Mj}(\sigma)]$ can be cut into three segments: one in the axis of ϕ^{Mj} and two "error" parts outside of the axis. Using the fact that (A) fails, we see that the part in the axis must not contain any long legal subsegment, thus its length can be estimated by counting the number of maximal legal subsegments in it. Lemma 3.4.2 controls the decrease of the number of maximal legal subsegments in σ up to error parts.

However the error parts may grow, for two reasons: they are stretched by f^M and the inner part produces small errors too, which add to the previous error. The aim of the proof is to take σ long enough to keep the growth of these error parts small in comparison with the decrease of the inner part, so that the overall effect of f^{Mi} on σ is a decrease.

Then we apply the reverse: we look at the evolution of $f_-^{Mj} \circ h \circ f^{Mi}(\sigma)$ for $j \in \{0, \dots, i\}$. Now the argument for the absence of long legal subsegments in the inner part is the fact that (B) fails, and the conclusion is similar: the overall length of the segment decreases. As a result $[f_-^{Mi} \circ h \circ f^{Mi}(\sigma)]$ is a lot shorter than σ in proportion.

The contradiction comes from Lemma 3.2.1: the maps h and $f_-^{Mi} \circ h \circ f^{Mi}$ are equal up to a bounded error, so when σ is long enough, it cannot decrease much in proportion when applying $f_-^{Mi} \circ h \circ f^{Mi}$.

We now write a formal argument along these lines.

Let $i \in \mathbb{N}$; set $N := Mi$. Let σ be a segment in $\mathrm{Axis}_T(g)$.

By Lemma 3.2.1, there exists a constant B_i such that for any $x \in T$, $d_{T_-}(h(x), f_-^{Mi} \circ h \circ f^{Mi}(x)) \leq B_i$. There exists a constant $L_{1,i}$ depending on B_i and h such that if $\mathrm{len}(\sigma) > L_{1,i}$

then $\mathrm{len}([h(\sigma)]) > 4B_i$ so

$$\frac{\mathrm{len}([f_-^{Mj} \circ h \circ f^{Mi}(\sigma)])}{\mathrm{len}(\sigma)} > \frac{1}{2} \tag{3.1}$$

Moreover there exists $D_i \geq 0$ depending on f, f_-, h, i such that for any $\sigma \subset \mathrm{Axis}_T(g)$, for any segment $\theta \subset \sigma$ at distance greater than D_i from the endpoints of σ, $[f_-^{Mi} \circ h \circ f^{Mi}(\theta)] \cap \mathrm{Axis}_T(g) \subset [h(\sigma)] \cap \mathrm{Axis}_T(g)$. Let $\tilde{\sigma}$ be the subsegment of σ obtained by cutting out a D_i-neighbourhood of the endpoints.

Suppose that both (A) and (B) fail for σ and for $N = Mi$.

Since (A) fails for σ, no segment $[f^{Mj}(\sigma)] \cap \mathrm{Axis}_T(\phi^{Mj}(g))$ for $j \in \{0, \ldots, i\}$ can contain a legal subsegment longer than C. By Lemma 3.4.2 there exist constants $K < 1, K' \geq 0$ and a decomposition $[f^M(\sigma)] = \alpha_1 \cdot \sigma_1 \cdot \beta_1$ such that

$$\mathrm{len}(\sigma_1) \leq K\,\mathrm{len}(\sigma) + K'$$

Define by induction $\alpha_j \cdot \sigma_j \cdot \beta_j = [f^M(\sigma_{j-1})]$ using Lemma 3.4.2. For all $j \in \{1, \ldots, i\}$ we have

$$\mathrm{len}(\sigma_j) \leq K\,\mathrm{len}(\sigma_{j-1}) + K'$$

and $\mathrm{len}(\alpha_j), \mathrm{len}(\beta_j) \leq K'/2$ so

$$\mathrm{len}([f^M(\sigma_{j-1})]) \leq K\,\mathrm{len}(\sigma_{j-1}) + K' + 2\lambda^M \frac{K'}{2}$$

where $\lambda := \mathrm{Lip}(f)$. Thus we have

$$\mathrm{len}([f^{Mi}(\sigma)]) \leq K^i\,\mathrm{len}(\sigma) + K'(1+\lambda^M)\sum_{j=0}^{i-1} K^j \leq K^i\,\mathrm{len}(\sigma) + \frac{K'(1+\lambda^M)}{1-K}$$

Remember that $\tilde{\sigma}$ be the subpath obtained from σ by cutting out the D_i-neighbourhood of the endpoints. Let $\sigma'_j := [f_-^{Mj} \circ h \circ f^{Mi}(\tilde{\sigma})] \cap \mathrm{Axis}_{T_-}(\phi^{M(i-j)}(g))$ for $j \in \{0, \ldots, i\}$: it cannot contain legal subsegments longer than C. Indeed $[f_-^{Mi} \circ h \circ f^{Mi}(\tilde{\sigma})] \cap \mathrm{Axis}_T(g) \subset [h(\sigma)] \cap \mathrm{Axis}_T(g)$ and by assumption that (B) fails, the latter does not contain any legal subsegments longer than C.

Applying the same argument as above using Lemma 3.4.2 with f_-, C, we obtain again

K_-, K'_- such that

$$\operatorname{len}([f_-^{Mi} \circ h \circ f^{Mi}(\tilde{\sigma})]) \leq K_-^i \operatorname{len}([h \circ f^{Mi}(\tilde{\sigma})]) + \frac{K'_-(1+\lambda_-^M)}{1-K_-}$$

with $\lambda_- := \operatorname{Lip}(f_-)$. By combining both inequalities, using the fact that h is Lipschitz:

$$\begin{aligned}
\operatorname{len}([f_-^{Mi} \circ h \circ f^{Mi}(\tilde{\sigma})]) &\leq \operatorname{Lip}(h) K_-^i \operatorname{len}([f^{Mi}(\tilde{\sigma})]) + \frac{K'_-(1+\lambda_-^M)}{1-K_-} \\
&\leq \operatorname{Lip}(h) K_-^i \operatorname{len}([f^{Mi}(\sigma)]) + 2\operatorname{Lip}(h) K_-^i D_i \operatorname{Lip}(f^M)^i + \frac{K'_-(1+\lambda_-^M)}{1-K_-} \\
&\leq \operatorname{Lip}(h)(K_-K)^i \operatorname{len}(\sigma) + \operatorname{Lip}(h) K_-^i \left(\frac{K'(1+\lambda^M)}{1-K} + 2D_i \operatorname{Lip}(f^M)^i \right) \\
&\qquad + \frac{K'_-}{1-K_-}
\end{aligned}$$

so

$$\operatorname{len}([f_-^{Mi} \circ h \circ f^{Mi}(\sigma)]) \leq \operatorname{Lip}(h)(K_-K)^i \operatorname{len}(\sigma) + S_i$$

where S_i is an additive constant depending on i.

Now assume that i is big enough so that $\operatorname{Lip}(h)(K_-K)^i < 1/4$. Note that the choice of i does not depend on σ but only on the maps f^M, f_-^M, h and on C. Then there exists a constant $L_{2,i} \geq 0$ such that if σ is longer than $L_{2,i}$ then

$$\frac{\operatorname{len}([f_-^{Mi} \circ h \circ f^{Mi}(\sigma)])}{\operatorname{len}(\sigma)} \leq \frac{1}{2} \tag{3.2}$$

If σ is longer than $L := \max\{L_{1,i}, L_{2,i}\}$ then inequations 3.1 and 3.2 contradict each other. This achieves the proof. □

Corollary 3.4.4. *For every $C > 0$ there exists $N \in \mathbb{N}$ such that for every $g \in G$ one of the followings holds:*

(A) $\operatorname{Axis}_T(\phi^N(g))$ *contains a legal segment of length $> C$*

(B) $\operatorname{Axis}_{T_-}(\phi^{-N}(g))$ *contains a legal segment of length $> C$.*

Proof. It suffices to apply Lemma 3.4.3 to a long enough subsegment $\sigma \subset \operatorname{Axis}_T(g)$. □

In the rest of the section, our aim will be to prove that if $g \in G$ then $\operatorname{Axis}_{T \cdot \phi^n(g)}(g)$ has increasingly long legal subsegments. This is a key step in the definition of the projection $\mathcal{D} \to \mathcal{L}_f$.

As in [Alg11] we define the legality threshold and the legality of a path in T:

Definition 3.4.5. Let $\kappa := \frac{4\,\mathrm{BBT}(f)}{\lambda-1} = 2C_f$ be the *legality threshold*.

For every finite path $\alpha \subset T$ we define the *legality ratio* of α with respect to the train track structure as follows. Let $\alpha_1, \dots, \alpha_k$ be the maximal legal subsegments of α. Then the legality of α is

$$\mathrm{Leg}_f(\alpha, T) := \frac{\sum_{\mathrm{len}(\alpha_i)\geq\kappa} \mathrm{len}(\alpha_i)}{\mathrm{len}(\alpha)}$$

which is the proportion of α which belongs to a legal subpath longer than κ.

If g is a loxodromic element of G, then we distinguish two cases:

— g is legal and we define $\mathrm{Leg}_f(g, T) = 1$

— there exists a fundamental domain α for g which starts and ends at an illegal turn of the axis. Then $\mathrm{Leg}_f(g, T) = \mathrm{Leg}_f(\alpha, T)$.

Remarks 3.4.6.

— If α is a fundamental domain of g which does not start and end at an illegal turn while $\mathrm{Axis}_T(g)$ contains one, then $\mathrm{Leg}_f(\alpha) \leq \mathrm{Leg}_f(g)$.

— If $(\alpha_n)_{n\in\mathbb{N}}$ is a sequence of nested subsegments of $\mathrm{Axis}_T(g)$ whose length goes to infinity then $\mathrm{Leg}_f(g, T) = \lim_{n\to\infty} \mathrm{Leg}_f(\alpha, T)$.

— For $l \in \mathbb{Z} \setminus \{0\}$ we have $\mathrm{Leg}_f(g^l, T) = \mathrm{Leg}_f(g, T)$.

— To define $\mathrm{Leg}_{f_-}(g, T_-)$ we use the threshold $\kappa_- := 2C_{f_-}$.

The following result states that if α contains sufficiently many long legal subsegments, then the length of $f^n(\alpha)$ grows exponentially, as though α were legal.

Lemma 3.4.7. *Let $\varepsilon > 0$. There exists a constant $C(\varepsilon)$ such that for every finite path α in T such that $\mathrm{Leg}_f(\alpha, T) \geq \varepsilon$, for every $n \in \mathbb{N}$ we have $\mathrm{len}([f^n(\alpha)]) \geq C(\varepsilon)\lambda^n \mathrm{len}(\alpha)$.*

Proof. Since $\mathrm{Leg}_f(\alpha, T) \geq \varepsilon$, α contains at least one legal subsegment of length greater than κ. Let $\beta_1, \dots, \beta_k$ be the maximal legal subsegments of α longer than κ. For every $i \in \{1, \dots, k\}$ let θ_i be the subsegment of β_i obtained by cutting out the $\frac{C_f}{2}$-neighbourhood of the endpoints. By Lemma 3.1.13 the images of $\theta_1, \dots, \theta_k$ by f^n are disjoint for any $n \in \mathbb{N}$ and contained in $[f^n(\alpha)]$. Moreover for every $i \in \{1, \dots, k\}$, $\mathrm{len}(\theta_i) \geq \frac{1}{2}\mathrm{len}(\beta_i)$.

Since $\kappa = 2C_f$ we have

$$\begin{aligned}\mathrm{len}([f^n(\alpha)]) &\geq \sum_{i=1}^{k} \mathrm{len}(f^n(\theta_i)) \\ &\geq \lambda^n \sum_{i=1}^{k} \mathrm{len}(\theta_i) \\ &\geq \frac{1}{2}\lambda^n \sum_{i=1}^{k} \mathrm{len}(\beta_i) \\ &\geq \frac{1}{2}\lambda^n \varepsilon\, \mathrm{len}(\alpha)\end{aligned}$$

Therefore we obtain the desired result, with $C(\varepsilon) = \varepsilon/2$. □

Corollary 3.4.8. *Let $\varepsilon > 0$. There exists a constant $C(\varepsilon)$ such that for every loxodromic $g \in G$ such that $\mathrm{Leg}_f(g, T) \geq \varepsilon$, for every $n \in \mathbb{N}$ we have $\|\phi^n(g)\|_T \geq C(\varepsilon)\lambda^n \|\phi^n(g)\|$.*

Proof. Let $\varepsilon > 0$. Let $g \in G$. There exists $x \in \mathrm{Axis}_T(g)$ such that $\mathrm{Leg}_f([x, gx], T) = \mathrm{Leg}_f(g, T) \geq \varepsilon$. By Lemma 3.4.7, for any $k \in \mathbb{N}$ we have $d_T(f^n(x), \phi^n(g^k)f^n(x)) \geq C(\varepsilon)\lambda^n d_T(x, g^k x)$.

Thus

$$\begin{aligned}\|\phi^n(g)\|_T &= \inf_{k\in\mathbb{N}} \frac{d_T(f^n(x), \phi^n(g^k)f^n(x)}{k} \\ &\geq C(\varepsilon)\lambda^n \frac{d_T(x, g^k x)}{k} \\ &= C(\varepsilon)\lambda^n \|g\|_T\end{aligned}$$

□

Lemmas 3.4.9 and 3.4.11 aim to prove basic properties which can be deduced from Lemma 3.4.7. Together they prove that the legality function $n \mapsto \mathrm{Leg}_f(\phi^n(g), T)$ cannot be greater than ε in a neighbourhood of $-\infty$. They are illustrated by Figure 3.8.

Lemma 3.4.9. *For any $\varepsilon > 0$ there exists $M \in \mathbb{N}$ such that for any $g \in G$ such that $\mathrm{Leg}_f(g) \geq \varepsilon$, for any $m \geq M$, $\|\phi^m(g)\|_T > \|g\|_T$.*

Proof. By applying Lemma 3.4.7 to a fundamental domain for g starting at a legal turn, there exists $C(\varepsilon)$ such that for all $m \in \mathbb{N}$

$$\|\phi^m(g)\|_T \geq C(\varepsilon)\lambda^m \|g\|_T$$

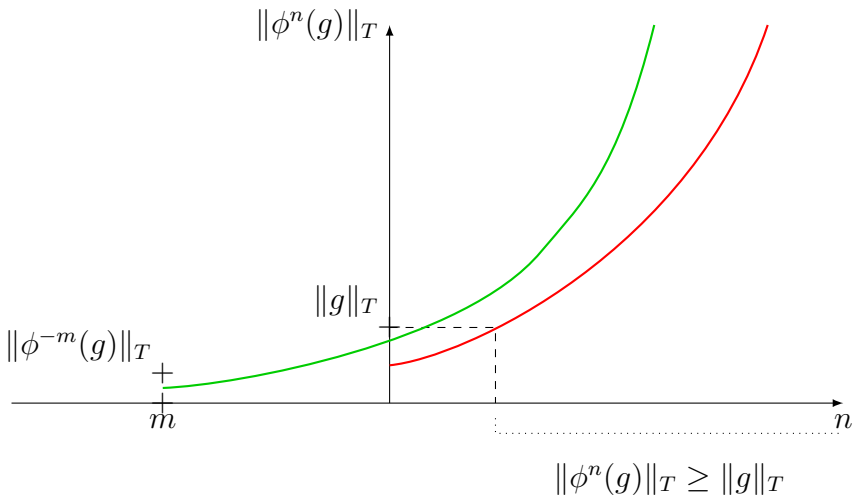

Figure 3.8 – If $\mathrm{Leg}_f(g,T) \geq \varepsilon_0$ then $\|\phi^n(g)\|_T$ is above the red graph. If $\mathrm{Leg}_f(\phi^{-m}(g),T) \geq \varepsilon_0$ then $\|g\|_T$ must be above the green graph, so if m is big then $\|\phi^{-m}(g)\|_T$ is small.

so with $M \geq -\frac{\log C}{\log \lambda}$ we get the lemma. □

Remark 3.4.10. Similarly there exists M_- such that if $\mathrm{Leg}_{f_-}(g) \geq \varepsilon$ then for any $m \geq M_-$, $\|\phi^{-m}(g)\|_{T_-} > \|g\|_{T_-}$.

Lemma 3.4.11. *For any loxodromic $g \in G$ there exists $m_g \in \mathbb{N}$ such that for any $m \geq m_g$,* $\mathrm{Leg}_f(\phi^{-m}(g)) < \varepsilon$.

Proof. Suppose $\mathrm{Leg}_f(\phi^{-m}(g)) \geq \varepsilon$ for some $m \in \mathbb{N}$ and $\varepsilon > 0$. By applying Corollary 3.4.8 to $\phi^{-m}(g)$, there exists $C(\varepsilon)$ such that $\|g\|_T \geq C(\varepsilon)\lambda^m \|\phi^{-m}(g)\|_T$. Let l_e be the length of the shortest edge in T, then $\|\phi^{-m}(g)\|_T \geq l_e$ so

$$m \leq \frac{\log(\|g\|_T) - \log C - \log l_e}{\log \lambda}.$$

□

Corollary 3.4.4 proves that for any g in G, either $\phi^N(g)$ has a f-legal segment of length C, either $\phi^{-N}(g)$ has an f_--legal segment of length C, where the integer N does not depend on g at all. A crucial point is Lemma 3.4.12, i.e. that such a result also works with the legality ratio, i.e. up to choosing a greater N, either the legality ratio of $\phi^N(g)$ in T or the legality ratio of $\phi^{-N}(g)$ in T_- is greater than a definite ε_0. Combined with Lemma 3.4.7 we will then be able to prove that $\mathrm{len}(g)$ grows exponentially when $n \to \pm\infty$, and has a minimum in a bounded subset of $\mathcal{L}_f$.

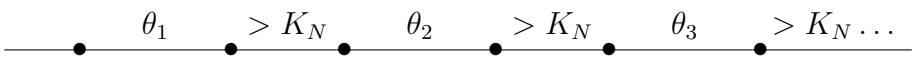

Figure 3.9 – The axis of g in T is cut into subsegments θ_i separated by subsegments longer than K_N.

The following lemma needs the fact that ϕ is pseudo-atoroidal since it relies on Lemma 3.1.18 through Corollary 3.4.4. It is proved for the free group case in [BFH97].

Lemma 3.4.12. *There exists $\varepsilon_0 > 0$ and $N \in \mathbb{N}$ such that for every loxodromic element $g \in G$, one of the followings holds:*

— $\mathrm{Leg}_f(\phi^N(g), T) > \varepsilon_0$

— $\mathrm{Leg}_{f_-}(\phi^{-N}(g), T_-) > \varepsilon_0$

Proof. Fix a G-equivariant quasi-isometry $h : T \to T_-$. Recall that for a G-equivariant quasi-isometry between G-trees $u : T_1 \to T_2$, such as h, f, f_- and their products, if a segment σ is contained in the axis of an element g in T_1, then $[u(\sigma)]$ is contained in the axis of g in T_2 apart from a $\mathrm{BBT}(u)$-neighbourhood of its endpoints.

Let $C := \max\{2C_f, 2C_{f_-}\}$. Let L, N be the constants given by Lemma 3.4.3.

There exists a constant K_N depending on L and the quasi-isometry constants for f, f_- such that for any points $x, y \in T$, $f^N(x) = f^N(y) \Rightarrow d_T(x, y) \leq K_N$ and $f_-^N \circ h(x) = f_-^N \circ h(y) \Rightarrow d_T(x, y) \leq K_N$).

Observe that for any subsegment $\sigma \subset \mathrm{Axis}_T(g)$ of length at least K_N, there exists $x \in \sigma$ such that $f^N(x) \in \mathrm{Axis}_T(\phi^N(g))$ and $x_- \in \sigma$ such that $f_-^N \circ h(x_-) \in \mathrm{Axis}_{T_-}(\phi^{-N}(g))$.

The axis of g in T can be cut into subsegments $\theta_i, i \in \mathbb{Z}$ of length L separated by other subsegments of length K_N (see Figure 3.9).

By the choice of K_N, for any $i \in \mathbb{Z}$, $[f^N(\theta_i)] \cap [f^N(\theta_{i+1})] = \varnothing$ and $[f_-^N \circ h(\theta_i)] \cap [f_-^N \circ h(\theta_{i+1})] = \varnothing$.

There exists a power g^l such that $2L + 3K_N < \|g^l\|_T$. Let $k := \lfloor \frac{\|g^l\|_T - K_N}{L + K_N} \rfloor$. There exists a fundamental domain α for g^l in $\mathrm{Axis}_T(g)$ which contains at least k consecutive $\theta_i, \theta_{i+1}, \ldots, \theta_{i+k-1}$ of the segments defined above, and at distance at least K_N from its endpoints.

Let $\alpha_N \subset T$ be a fundamental domain for $\phi^N(g^l)$ contained in $[f^N(\alpha)]$, and let $\alpha_{-N} \subset T_-$ be a fundamental domain for $\phi^{-N}(g^l)$ contained in $[f_-^N \circ h(\alpha)]$. Since there is a K_N-margin between θ_i, θ_{i+k-1} and the endpoints of α, for every $j \in \{i, \ldots, i + k - 1\}$ we have $[f^N(\theta_j)] \cap \mathrm{Axis}_T(\phi^N(g)) \subset \alpha_N$ and $[f_-^N(\theta_j)] \cap \mathrm{Axis}_{T_-}(\phi^{-N}(g)) \subset \alpha_{-N}$.

By Lemma 3.4.3, for each $j \in \mathbb{Z}$, either $[f^N(\theta_j)] \cap \mathrm{Axis}_T(\phi^N(g))$ contains an f-legal segment with length C, or $[f_-^N \circ h(\theta_j)] \cap \mathrm{Axis}_{T_-}(\phi^{-N}(g))$ contains an f_--legal segment with length C. Suppose the first case happens for at least half of the indices in $\{i, \dots, i+k-1\}$. Then since the images of the segments θ_j do not overlap, there are at least $k/2$ legal segments with length C in the fundamental domain α_N. Thus $\mathrm{Leg}_f(\phi^N(g)) \geq \frac{Ck}{2\,\mathrm{len}(\alpha_N)}$.

Since $k > \frac{\mathrm{len}(\alpha) - L - 2K_N}{L + K_N}$ and the fact that $\|g^l\|_T = \mathrm{len}(\alpha) \geq 2L + 3K_N$ we obtain

$$\begin{aligned}
\mathrm{Leg}_f(\phi^N(g)) &\geq \frac{C}{2\,\mathrm{len}(\alpha_N)} \frac{\mathrm{len}(\alpha) - L - 2K_N}{L + K_N} \\
&\geq \frac{C\,\mathrm{len}(\alpha)}{2\,\mathrm{len}(\alpha_N)} \frac{1 - \frac{L-2K_N}{2L+3K_N}}{L + K_N} \\
&\geq \frac{C\,\mathrm{len}(\alpha)}{2\,\mathrm{len}(\alpha_N)} \frac{L + 5K_N}{(L + K_N)(2L + 3K_N)} \\
&\geq \frac{C}{2\,\mathrm{Lip}(f^N)} \frac{L + 5K_N}{(L + K_N)(2L + 3K_N)} > 0
\end{aligned}$$

The bound does not depend on g nor on α.

Similarly, if the second case happens, i.e. if there are more long legal segments in $\alpha_{-N} \subset \mathrm{Axis}_{T_-}(\phi^{-N}(g))$, then we get

$$\mathrm{Leg}_{f_-}(\phi^{-N}(g)) \geq \frac{C}{2\,\mathrm{Lip}(f_-^N \circ h)} \frac{1 - \frac{L-2K_N}{2L+3K_N}}{L + K_N}$$

Since at least one of these two cases occurs, we can define ε_0 as the smallest of both bounds and we obtain the lemma. □

For a geodesic in T (resp. T_-) and a constant $L > 0$ we define the *lamination ratio* $\mathrm{LR}(g, T, \Lambda^+, L)$ (resp. $\mathrm{LR}(g, T_-, \Lambda^-, L)$ as the upper bound of the proportion of $\mathrm{Axis}_T(g)$ (resp. $\mathrm{Axis}_{T_-}(g)$) which can be covered by pairwise disjoint leaf segments of Λ^+ (resp. Λ^-) with length at least L.

Lemma 3.4.13. *Let $\varepsilon_0 > 0$. For any $L > 0$, there exists $N_1 \geq 0$ such that for any $g \in G$, for any $n \geq N_1$, if $\mathrm{Leg}_f(g, T) > \varepsilon_0$ then*

$$\mathrm{LR}(\phi^n(g), T, \Lambda^+, L) > \varepsilon_0/4.$$

Proof. Let $g \in G$ be such that $\mathrm{Leg}_f(g, T) > \varepsilon_0$. Let $\beta \subset \mathrm{Axis}_T(g)$ be a maximal legal subsegment with length at least κ. Let $\beta' \subset \beta$ be the subsegment obtained by cutting out

the $\frac{C_f}{2} \leq \frac{\kappa}{4}$-neighbourhood of the endpoints. Its length is at least $\text{len}(\beta) - C_f$ and for all $n \in \mathbb{N}$, $f^n(\beta') \subset \text{Axis}_T(\phi^n(g))$.

There exists $n_1 \in \mathbb{N}$ such that $\lambda^{n_1} C_f \geq 4l_{\max}$ where $l_{\max} := \max_{e \in E(T)} \text{len}(e)$. Thus $f^{n_1}(\beta')$ contains at least one edge of T. In fact, the number of edges of T contained in $f^{n_1}(\beta')$ is at least $k_{\beta'} := \left\lfloor \frac{\lambda^{n_1}(\text{len}(\beta) - C_f)}{l_{\max}} \right\rfloor$ and their total length is at least $c_{\beta'} := \lambda^{n_1}(\text{len}(\beta) - C_f) - 2l_{\max}$.

There exists $n_2 \in \mathbb{N}$ such that for every edge $e \in T$, $f^{n_2}(e)$ is a leaf segment with length greater than L.

Thus $f^{n_1+n_2}(\beta')$ is contained in $\text{Axis}_T(\phi^{n_1+n_2}(g))$ and contains at least $k_{\beta'}$ disjoint open leaf segment with length at least L whose total length is at least $\lambda^{n_2} c_{\beta'}$.

Let $\mathcal{B}$ be the set of $\langle g \rangle$-orbits of maximal legal subsegments of $\text{Axis}_T(g)$.

Let $n \in \mathbb{N}$. The proportion of $\text{Axis}_T(\phi^{n_1+n_2+n}(g))$ covered by the leaf segments is at least

$$\frac{\lambda^{n_2+n} \sum_{\beta \in \mathcal{B}} \lambda^{n_1}(\text{len}(\beta) - C_f) - 2l_{\max}}{\lambda^{n_1+n_2+n} \|g\|_T}$$

Thus

$$\begin{aligned}
\text{LR}(\phi^{n_1+n_2+n}(g), T, \Lambda^+, L) &\geq \frac{\sum_{\beta \in \mathcal{B}} \lambda^{n_1}(\text{len}(\beta) - C_f) - 2l_{\max}}{\lambda^{n_1} \|g\|_T} \\
&\geq \frac{\sum_{\beta \in \mathcal{B}} \lambda^{n_1}(\text{len}(\beta) - C_f) - \lambda^{n_1} \frac{C_f}{2}}{\lambda^{n_1} \|g\|_T} \\
&\geq \frac{\sum_{\beta \in \mathcal{B}} \frac{\text{len}(\beta)}{4}}{\|g\|_T} \\
&\geq \frac{\varepsilon_0}{4}
\end{aligned}$$

This proves the lemma with $N_1 = n_1 + n_2$. □

Lemma 3.4.12 yields a constant ε_0. For the rest of the paper we fix such an ε_0. Define $k(g), k_-(g)$ as follows:

— $k(g) = \min\{k \in \mathbb{Z} / \text{Leg}_f(\phi^k(g)) \geq \varepsilon_0\}$

— $k_-(g) = \max\{k \in \mathbb{Z} / \text{Leg}_{f_-}(\phi^k(g)) \geq \varepsilon_0\}$

By Lemma 3.4.11 these integers are well-defined.

Lemma 3.4.14. *There exists $N \in \mathbb{N}$ such that for any loxodromic $g \in G$, $|k(g) - k_-(g)| \leq N$.*

Proof. Let $g \in G$ be a loxodromic element. Let N_0 be the constant given by Lemma 3.4.12. Lemma 3.4.12 implies that either $\mathrm{Leg}_f(\phi^{N_0}(\phi^{k(g)-N_0-1}(g)))$ or $\mathrm{Leg}_{f_-}(\phi^{-N_0} \circ \phi^{k(g)-N_0-1}(g))$ is greater than ε_0. By definition of $k(g)$, the former does not hold so $\mathrm{Leg}_{f_-}(\phi^{-2N_0+k(g)-1}(g)) > \varepsilon_0$. Therefore we have

$$k(g) - k_-(g) \leq 2N_0 + 1.$$

Lemma 3.4.9 gives M such that for all $m \geq M$ we have

$$\|\phi^m(\phi^{k(g)}(g))\|_T > \|\phi^{k(g)}(g)\|_T.$$

It also gives a similar constant M_- for f_-. If we had $k_-(g) - k(g) \geq \max\{M, M_-\}$ then by applying Lemma 3.4.9 twice with $m = k_-(g) - k(g)$ we would get a contradiction:

$$\|\phi^k(g)\|_T = \|\phi^{-m} \circ \phi^m \circ \phi^k(g)\|_T > \|\phi^m \circ \phi^k(g)\|_T > \|\phi^k(g)\|_T$$

This gives an upper bound for $k_-(g) - k(g)$. □

3.5 Defining the projection

Let $g \in G$ be a loxodromic element. Like in [Alg11] we define $t_0(g) := k(g)\log(\lambda_+)$. The following lemma is the same as [Alg11, Lemma XX]. The fact that G is a GBS group instead of F_N has no influence.

Lemma 3.5.1. *There exists a constant $C > 0$ such that for every loxodromic element $g \in G$ we have for $t \geq t_0$:*

$$C^{-1}\lambda^{\left\lfloor \frac{t-t_0}{\log(\lambda)} \right\rfloor}\|g\|_{T_0} \leq \|g\|_{T_t} \leq C\lambda^{\left\lfloor \frac{t-t_0}{\log(\lambda)} \right\rfloor}\|g\|_{T_0}$$

and for $t \leq t_0$:

$$C^{-1}\lambda_-^{\left\lfloor \frac{t_0-t}{\log(\lambda)} \right\rfloor}\|g\|_{T_0} \leq \|g\|_{T_t} \leq C\lambda_-^{\left\lfloor \frac{t_0-t}{\log(\lambda)} \right\rfloor}\|g\|_{T_0}$$

Proof. We will prove the inequalities in the case where t is a multiple of $\log(\lambda)$. The result for other values of t can be obtained by applying Lemma 3.1.27 to a translate of the subsegment $\{T_t / 0 \leq t \leq \log(\lambda)\}$, and it will only result in increasing the multiplicative constants by a controlled amount.

Write $t_0 = t_0(g)$. First let us deal with the case $t \geq t_0$. Let $n \in \mathbb{N}$ and let $t_n = t_0 + n \log \lambda$.

We have $T_{t_n} = T_{t_0} \cdot \phi^n$.

Since f is λ-Lipschitz we have

$$\|g\|_{T_{t_n}} \leq \lambda^n \|g\|_{T_{t_0}}$$

Let us prove the other side of the inequality. Lemma 3.4.7 can be applied to a well-chosen fundamental domain for g and gives a constant $C(\varepsilon_0)$ independant of n and g such that

$$\|g\|_{T_{t_n}} \geq C(\varepsilon_0)\lambda^n \|g\|_{T_{t_0}}$$

which gives the first part of the Lemma.

Now let us deal with the case $t \leq t_0$. Let $n \in \mathbb{N}$. Let $t_n = t_0 - n \log(\lambda)$. In that case we have $n = \frac{t_0 - t}{\log(\lambda)}$.

We have

$$\mathrm{Lip}(T, T_-)^{-1} \leq \frac{\|g\|_T}{\|g\|_{T_-}} \leq \mathrm{Lip}(T_-, T)$$

and by applying this to $\phi^n(g)$ instead of g for any $n \in \mathbb{Z}$, we have

$$\mathrm{Lip}(T, T_-)^{-1} \leq \frac{\|g\|_{T \cdot \phi^n}}{\|g\|_{T_- \cdot \phi^n}} \leq \mathrm{Lip}(T_-, T)$$

In particular, since $T_{t_0} = T \cdot \phi^{k(g)}$ this also works when replacing T, T_- with T_{t_0}, T_{-t_0}.

We now have $T_{-t_n} = T_{-t_0} \cdot \phi^{-n}$ so $\|g\|_{T_{-t_n}} \leq \lambda_-^n \|g\|_{T_{-t_0}}$. We deduce the right inequality:

$$\begin{aligned}
\|g\|_{T_{t_n}} &\leq \mathrm{Lip}(T_-, T)\|g\|_{T_{-t_n}} \\
&\leq \mathrm{Lip}(T_-, T)\lambda_-^n \|g\|_{T_{-t_0}} \\
&\leq \mathrm{Lip}(T_-, T)\,\mathrm{Lip}(T, T_-)\lambda_-^n \|g\|_{T_{t_0}}
\end{aligned}$$

Now let us prove the left inequality. Lemma 3.4.12 gives an integer constant N such that $k_-(g) \geq k(g) - N$. Thus we have by the same arguments as above, and for $n \geq N$ and $T_{t_N} = T_{t_0} \cdot \phi^{-N}$ we obtain

$$\|g\|_{T_{-t_n}} \geq C(\varepsilon_0)\lambda_-^{n-N} \|g\|_{T_{-t_N}}$$

Since $\|g\|_{T_{-t_N}} \geq \mathrm{Lip}(T_- \cdot \phi^{-N}, T_-)^{-1}\|g\|_{T_{-t_0}}$ we have

$$\|g\|_{T_{-t_n}} \geq C(\varepsilon_0)\lambda_-^{-N}\,\mathrm{Lip}(T_- \cdot \phi^{-N}, T_-)^{-1}\lambda_-^n\|g\|_{T_{-t_0}}$$

When $0 \leq n < N$ then we have $\frac{\|g\|_{T_{-t_0}}}{\|g\|_{T_{-t_n}}} \leq \mathrm{Lip}(T_-, T_- \cdot \phi^N)$ so in any case there is a constant $C > 1$ depending only on T, T_-, N such that

$$\|g\|_{T_{t_n}} \geq C^{-1}\lambda_- n(t)\|g\|_{T_{t_0}}$$

This proves the lemma. □

Define $\Theta(g) := \{t \in \mathbb{R} / \|g\|_{T_t} \text{ minimal }\}$.

Lemma 3.5.2. *There exists $s > 0$ such that for every loxodromic $g \in G$ and $t \in \Theta(g)$, then $|t - t_0(g)| < s$.*

Proof. Let C be the constant from Lemma 3.5.1. Suppose $t > t_0(g)$. Then we have $\|g\|_{T_t} \geq C^{-1}\lambda^{\left\lfloor \frac{t-t_0(g)}{\log(\lambda)} \right\rfloor}\|g\|_{T_{t_0}}$ so since $\|g\|_{T_t} \leq \|g\|_{T_{t_0}}$ this implies $t - t_0 \leq \log(C\lambda)$.

We get a similar inequality for $t < t_0$, hence the result. □

Remark 3.5.3. The diameter of $\Theta(g)$ is bounded by $2s$.

An important property of the projection is that projections of simple pairs are close:

Lemma 3.5.4. *There exists $s' > 0$ with the following property. Let $\{g, h\}$ be a simple pair of loxodromic elements of G. Then $|t_0(g) - t_0(h)| < s'$.*

Proof. We prove this by contraposition: we will show that if $t_0(g)$ and $t_0(h)$ are too far apart then we can find t in between such that the axes of g and h in T_t contain long segments of the stable and unstable lamination (see Figure 3.10).

Let L_0 be the constant from Proposition 3.3.13 such that elements in a simple pair cannot contain leaf segments longer than L_0 of both lamination in their axes. Without loss of generality we may assume $L_0 > \kappa$.

By Lemma 3.2.9 there exists a constant L_1 such that if a path $\beta \subset T_-$ contains a leaf segment of Λ^- longer than L_1, then $[h_-(\beta)]$ contains a leaf segment of Λ^- of length greater than L_0. The choice of L_1 depends only on h_-, f_-, and L.

Without loss of generality suppose $t_0(g) < t_0(h)$.

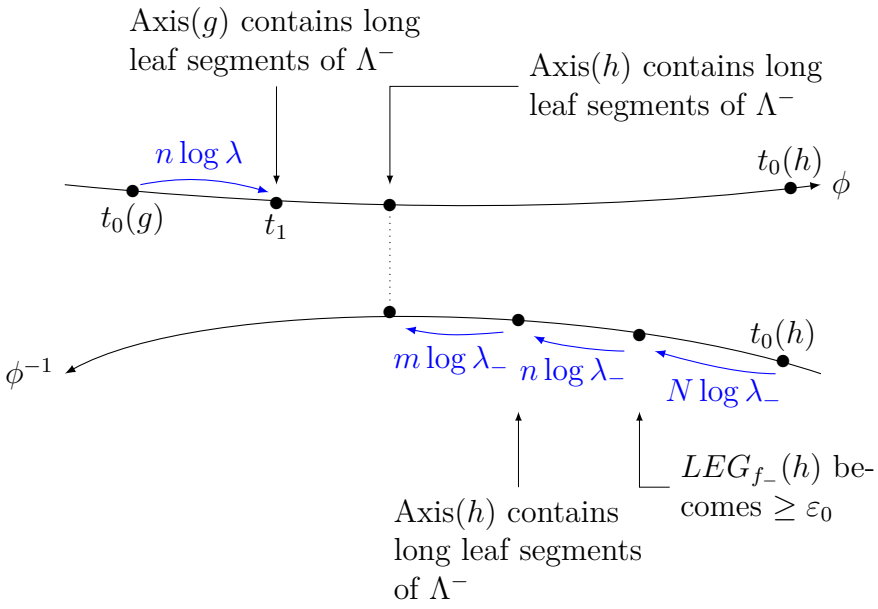

Figure 3.10 – If $t_0(g), t_0(h)$ are sufficiently far apart, then there exists t such that the axes of g and h in T_t contain long leaf segments of the opposite laminations

Let $N_1 > 0$ be the integer given by Lemma 3.4.13 for f, ε_0 and L_0. Similarly define $N_{1,-} > 0$ as the integer given for f_-, ε_0, L_1.

By definition of t_0 we have $\mathrm{Leg}_f(g, T_{t_0(g)}) \geq \varepsilon_0$. Thus for all $n \geq N_1$ we have

$$\mathrm{LR}(g, T_{t_0(g)+n\log(\lambda)}, \Lambda^+, L_0) \geq \varepsilon_0/4.$$

Similarly for all $n \geq N_{1,-}$ we have $\mathrm{Leg}_{f_-}(h, T_{t_0(h)-n\log(\lambda_-)}, \Lambda^-, L_1) \geq \varepsilon_0/4$.

Suppose $t_0(h) - t_0(g) > N_1 \log(\lambda) + N_{1,-} \log(\lambda_-)$. Then there exists $t_0(g) + N_1 \log(\lambda) < t < t_0(h) - N_{1,-} \log(\lambda_-)$. Consequently $\mathrm{Axis}_{T_t}(g)$ contains an L_0-piece of Λ^+ and $[h(\mathrm{Axis}_{T_{-t}}(h))]$ contains an L_0-piece of Λ^-. This contradicts the fact that the pair $\{g, h\}$ is simple. □

Here is a direct corollary:

Corollary 3.5.5. *Let s, s' be the constants from Lemmas 3.5.2 and 3.5.4. For a simple pair $\{g, h\}$, $\mathrm{diam}(\Theta(g) \cup \Theta(h)) < s + s'$.*

In order to evaluate the distance $d_{\mathrm{Lip}}(X, \mathcal{L}_f)$ for some arbitrary $X \in \mathcal{D}$, we will use candidates of X. Lemma 3.5.4 applies in particular to candidates:

Corollary 3.5.6. *Suppose $b_1(G) \geq 3$. There exists $s'' > 0$ such that for every $X \in \mathcal{D}$, if g, h are candidates in X, then for any $t_g \in \Theta(g)$ and $t_h \in \pi_f(h)$ we have $|t_g - t_h| < s''$.*

Proof. By Lemma 3.3.3 there exists $k \in G$ such that the pairs $\{g, k\}$ and $\{h, k\}$ are simple. Applying the previous corollary gives $|t_g - t_h| < |t_g - t_k| + |t_k - t_h| < 2s' + 2s$. □

For $X \in \mathcal{D}$, define $\Theta_X := \{t \in \mathbb{R}/d_{\text{Lip}}(X, T_t) \text{ minimal }\}$.

Lemma 3.5.7. *For every $X \in \mathcal{D}$, the set Θ_X is non-empty.*

Moreover there exists $s > 0$ such that for every $X \in \mathcal{D}$, $\text{diam}(\Theta_X) < s$.

Proof. Let $X \in \mathcal{D}$. By Theorem 3.1.23, for all $t \in \mathbb{R}$, there exists a candidate such that $d_{\text{Lip}}(X, T_t) = \log \frac{\|g\|_{T_t}}{\|g\|_X}$.

Therefore we have

$$d_{\text{Lip}}(X, T_t) = \max_{g \text{ candidate}} \log \frac{\|g\|_{T_t}}{\|g\|_X}$$

Fix a candidate g. The function $t \mapsto \frac{\|g\|_{T_t}}{\|g\|_X}$ is minimal for $t \in \Theta_g$. We will prove that $d_{\text{Lip}}(X, T_t)$ reaches its minimum in a D-neighbourhood of Θ_g, where D does not depend on X nor on the number of candidates in X.

Let $t_0 := t_0(g)$. If h is another candidate we have $|t_0(h) - t_0| < s''$ where s'' is the constant from Corollary 3.5.6. By Lemma 3.5.1 we have

$$\|h\|_{T_{t_0}} \leq \|h\|_{T_{t_0(h)}} C\lambda^{\frac{s''}{\log(\lambda)}}$$

Write $K = C\lambda^{\frac{s''}{\log(\lambda)}}$. For $t_* \geq \log(2CK) + \log(\lambda)$ and $t > t_0 + s'' + t_*$ we have $t > t_0(h) + t_*$ and still by Lemma 3.5.1 we get

$$\|h\|_{T_t} \geq 2K\|h\|_{T_{t_0(h)}} \geq 2\|h\|_{T_{t_0}}$$

Dividing both sides by $\|h\|_X$ does not change the inequality.

Therefore we have for $t > t_0 + s'' + t_*$

$$\max_{h \text{ candidate}} \frac{\|h\|_{T_t}}{\|h\|_X} \geq 2\frac{\|h\|_{T_{t_0}}}{\|h\|_X}$$

For $t < t_0$ we get a similar result. We deduce a constant Δt such that for $t \in \mathbb{R} \setminus [t_0 - \Delta t, t_0 + \Delta t]$ we have $d_{\text{Lip}}(X, T_t) > \log 2 d_{\text{Lip}}(X, T_{t_0})$. Since $t \to T_t$ is continuous for the axes topology, $t \mapsto d_{\text{Lip}}(X, T_t)$ reaches its minimum in a Δt-neighbourhood of t_0. □

Remark 3.5.8. The previous proof differs slightly from the proof of [Alg11, p. XX] since there is no bound on the number of candidates in elements of $\mathcal{D}$, unlike in CV_N. This

comes from the fact that $\mathcal{D}$ is not finite dimensional so there is no bound on the number of orbits of edges in elements of $\mathcal{D}$.

For $X \in \mathcal{D}$ we choose t_X in Θ_X. Since Θ_X has bounded diameter and the bound does not depend on X this will be well enough defined.

3.6 Negative curvature properties of the projection

In this section we prove the analogues of Lemmas 5.7 and 5.8 of [Alg11] and deduce the strong contraction property.

The difference in the proof with [Alg11] is the proof of Lemma 3.6.4. The initial proof relies on special shapes of graphs such as roses (see [Alg11, Proposition 5.10]). Here reduced graphs will take the role of roses. The other proofs are actually quite similar to the free group case.

Lemma 3.6.1. *There exist $s, c > 0$ such that for any $X \in \mathcal{D}$, if $|t - t_X| > s$ then $d_{\mathrm{Lip}}(X, T_t) \geq d_{\mathrm{Lip}}(X, \pi(X)) + d_{\mathrm{Lip}}(\pi(X), T_t) - c$.*

Proof. Suppose $t \leq t_X$. Let g be a candidate in X. The idea of the proof is that if s is big enough, then $\mathrm{Leg}_f(g, T_{t_X+s})$ is also big and g almost realizes $d_{\mathrm{Lip}}(T_{t_X+s}, T_t)$.

There is a candidate h of X such that $\mathrm{Lip}(X, \pi(X)) = \frac{\|h\|_{\pi(X)}}{\|h\|_X}$. Since $t_X \in \Theta(X)$ we have $t_X \in \Theta(h)$. By Lemmas 3.5.2 and 3.5.4 there exists a constant s such that for every candidate g of X we have $|t_X - t_0(g)| < s$. Thus for any candidate g of X, for any $t_1 > t_X + s$, we have $\mathrm{Leg}_f(g, T_1) > \varepsilon_0$.

Let $Z := T_{t_1}$. Let g be a candidate of X such that $d_{\mathrm{Lip}}(X, Z) = \frac{\|g\|_Z}{\|g\|_X}$.

Applying twice Lemma 3.5.1 to g for t and t_1 we obtain a constant C such that

$$\frac{\|g\|_{T_t}}{\|g\|_Z} \geq C^{-2}\lambda^{\frac{t-t_0(g)}{\log(\lambda)} - \frac{t_1-t_0(g)}{\log(\lambda)} - 1} = C^{-2}\lambda^{-1}e^{t-t_1}$$

Remarking that $e^{t-t_1} = \mathrm{Lip}(Z, T_t)$ we have

$$\mathrm{Lip}(Z, T_t) \leq \frac{\|g\|_{T_t}}{\|g\|_Z} \frac{1}{C^2\lambda}$$

with $C^2\lambda > 1$.

Since $\mathrm{Lip}(X,Z) = \frac{\|g\|_Z}{\|g\|_X}$ we have

$$\begin{aligned}\mathrm{Lip}(X,T_t) &\geq \frac{\|g\|_{T_t}}{\|g\|_X} \\ &= \frac{\|g\|_{T_t}}{\|g\|_Z}\frac{\|g\|_Z}{\|g\|_X} \\ &\geq \frac{1}{C^2\lambda}\mathrm{Lip}(Z,T_t)\,\mathrm{Lip}(X,Z)\end{aligned}$$

Applying the logarithm we get a constant $K > 0$ such that $d_{\mathrm{Lip}}(X,T_t) \geq d_{\mathrm{Lip}}(X,Z) + d_{\mathrm{Lip}}(Z,T_t) - K$.

Finally by definition of the projection we have $d_{\mathrm{Lip}}(X,Z) \geq d_{\mathrm{Lip}}(X,\pi(X)$. If $t - t_X > s$ we have

$$\begin{aligned}d_{\mathrm{Lip}}(X,T_t) &\geq d_{\mathrm{Lip}}(X,\pi(X) + d_{\mathrm{Lip}}(Z,T_t) - K \\ &\geq d_{\mathrm{Lip}}(X,\pi(X)) + d_{\mathrm{Lip}}(\pi(X),T_t) - s - K\end{aligned}$$

□

Lemma 3.6.2. *There exist $s, c > 0$ such that for any $X, Y \in \mathcal{D}$, if $|t_X - t_Y| > s$ then $d_{\mathrm{Lip}}(Y,X) \geq d_{\mathrm{Lip}}(Y,\pi(X)) - c$.*

Before proving Lemma 3.6.2 we need some preliminary results.

Lemma 3.6.3. *Let $X, T \in \mathcal{D}$ and $e_0 \in E(X)$. Suppose every edge orbit in $X \setminus G \cdot e_0$ is non-collapsible. There exists a G-equivariant map $\tau : X \to T$ such that every edge in $X \setminus G \cdot e_0$ is contained in a τ-legal bi-infinite geodesic in $X \setminus G \cdot e_0$.*

Proof. We will prove this by constructing the map $\tau : X \to T$ such that at every vertex $v \in V(X)$, at least two gates at v for the gate structure induced by τ contain edges in $E(X) \setminus G \cdot e_0$. Then there exist bi-infinite τ-legal geodesics with the desired property.

There exists a G-equivariant map $\tau_0 : X \to T$. We may suppose that τ_0 sends vertex to vertex and is linear on edges.

Let $v_1, \ldots, v_n$ be representatives of every vertex orbit of $V(X)$. In order to define a new map τ, it suffices to choose the image of v_i for every $i \in \{1, \ldots, n\}$. The image of v_i can be any vertex $w_i \in V(T)$ such that $G_{v_i} \subset G_{w_i}$.

Suppose there exists $v \in X$ such that there is only one gate at v which contains edges in $E(X) \setminus G \cdot e_0$. We include the case where there is only one gate at v.

Let $w = \tau_0(v)$. There exists a vertex $w' \in T$ such that $\bigcap_{e \in E_v \setminus G \cdot e_0} \tau_0(e) = [w, w']$. Since no edges in τ_0 are collapsible except translates of e_0, the images $\tau_0(e)$ have non-zero length. Since these edges are contained in a single gate, the intersection has non-zero length and we have $w \neq w'$.

Let us prove $G_v \subset G_{w'}$. Let $a \in G_v$. By contradiction suppose $aw' \neq w'$, then for any edge in $E_v \setminus G \cdot e_0$ we have $w' \notin \tau(ae)$ which contradicts the definition of w', thus $G_v \subset G_{w'}$. Define τ_1 by

$$\tau_1 : x \in V(X) \longmapsto \begin{cases} \tau_0(x) & \text{if } x \notin G \cdot v \\ gw' & \text{if } x = gv \end{cases}$$

Note that if there exists $e \in E_v$ such that $\tau_0(e) = [w, w']$, then e cannot be a loop in the quotient. If e were a loop $[v, gv] \subset T$ then $g^{-1}\bar{e} = [v, g^{-1}v]$, which is also in E_v and has same length, would also be sent to $[w, w'] = [w, gw] = [w, g^{-1}w]$. This would imply that g^2 is elliptic, which is a contradiction. Thus the vertices of e are in distinct orbits, and since the image of e is a single point, τ_1 factors through the collapse of e.

Since no edge orbit is collapsible except e_0, e can be collapsed only if e is a translate of e_0. We do not care about the image of e_0.

By construction of τ_1, there are at least two gates for τ_1 at vertex v which contain edges in $E_v \setminus G \cdot e_0$. Let us prove that gates at other vertices have not changed.

Let $e \in E(T) \setminus G \cdot e_0$. Neither τ_0 nor τ_1 collapse e. If an endpoint x of e is in $G \cdot v$, then $\tau_1(x)$ is in the interior of $\tau_0(e)$. Otherwise $\tau_1(x) = \tau_0(x)$. Thus if $o(e) \notin G \cdot v$, the first edge of $\tau_1(e)$ and $\tau_0(e)$ are equal. Thus the gate structure at $o(e)$ is unchanged.

Therefore the number of vertices with at most one gate containing edges not in $G \cdot e_0$ is smaller for τ_1 than for τ_0.

We can iterate this procedure with τ_1 instead of τ_0 until we find a map τ such that all vertices in X have at least two gates containing edges not in $G \cdot e_0$. □

Lemma 3.6.4. *Let L be the constant from Proposition 3.3.13. Suppose $r \in \mathbb{R}$ is such that for every candidate $u \in G$ for X, the axis $\mathrm{Axis}_{T_r}(u)$ contains an L-piece of Λ^-. Let $g \in G$. Suppose that $\mathrm{Axis}_{T_r}(g)/\langle g \rangle$ contains k disjoint $2L$-pieces of Λ^+ for some $k \in \mathbb{N}$. Then for every edge $e_0 \in E(X)$, $\mathrm{Axis}_X(g)/\langle g \rangle$ contains at least $k/2$ edges in the orbit of e_0.*

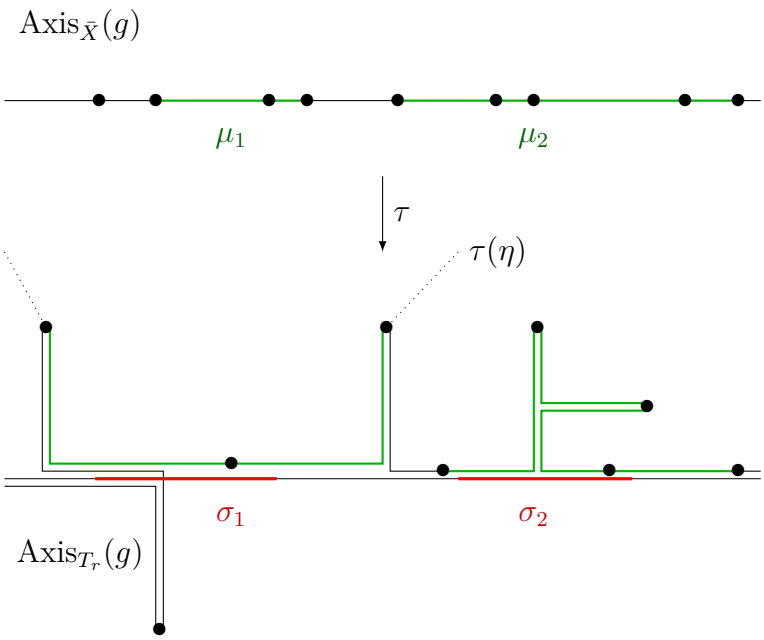

Figure 3.11 – Axis of g in T_r; the $2L$-pieces are in thick red, the image of $\mathrm{Axis}_{X'}(g)$ in dotted line and the μ_j are highlighted in green. The line $\tau(\eta)$ is also represented. Due to minimality of μ_j, η contains σ_j.

Proof. Let e_0 be an orbit of edges in $E(X)$. There exists a collapse $X \to X'$ such that e_0 is not sent to a single point, and every edge $e \neq e_0$ in X/G is not collapsible.

By Lemma 3.6.3 there exists a map $\tau : X' \to T_r$ such that:

— at every vertex $v \in V(X')$, there are at least two gates for the train track structure induced by τ

— at every vertex v, at least two gates contain edges which are not in $G \cdot e_0$.

Let $\sigma_1, \dots, \sigma_k$ be the $2L$-pieces of Λ^+ in $\mathrm{Axis}_{T_r}(g)$. Let $\mu_j := e_i \dots e_l$ be a minimal edge path in $\mathrm{Axis}_{X'}(g)$ such that $[\tau(\mu_j)]$ contains σ_j (see Figure 3.11). We claim that for all $j \in \{1, \dots, k\}$, the edge path μ_j contains one edge in $G \cdot e_0$. By contradiction, assume that for some $j \in \{1, \dots, k\}$ no edge in μ_j is in $G \cdot e_0$.

The path μ_j can be completed into a bi-infinite line $\eta \subset X' \setminus G \cdot e_0$ such that every turn of η is τ-legal, apart from turns in the interior of $e_i \dots e_l$. By minimality of μ_j, the image $\tau(e_i)$ (resp. $\tau(e_l)$) is not contained in $\tau(e_{i+1} \dots e_l)$ (resp. $\tau(e_i \dots e_{l-1})$). Therefore the legality property of η implies that the image $[\tau(\eta)]$ contains the segment $[\tau(\mu_j)]$.

Now we would like to find the axis of an element $h \in G$ such that $\mathrm{Axis}_{T_r}(h)$ contains

$[\tau(\mu_j)]$ and $\mathrm{Axis}_{X'}(h) \cap G \cdot e_0 = \varnothing$. Suppose we find such an h. Then $\mathrm{Axis}_{T_r}(h)$ contains an L-piece of Λ^+. However, in X', there exists a candidate $u \in G$, possibly equal to h, whose axis does not cross $G \cdot e_0$. The assumptions of the lemma imply that $\mathrm{Axis}_{T_r}(u)$ contains an L-piece of Λ^-. By Proposition 3.3.13 the pair $\{h, u\}$ is not simple. This contradicts the fact that their axes in X' both avoid $G \cdot e_0$. The conclusion is that e_0 must appear somewhere in μ_j.

Let us explain how we construct h. In the special case where there exists an edge $e \in E(X')$ such that there are two translates e, he in η with same orientation, and one on each side of μ_j then $[e, he]$ contains a fundamental domain for the axis of h and again by minimality of μ_j, $\mathrm{Axis}_{T_r}(h)$ contains $[\tau(\mu_j)]$.

In the general case, since no edge in $X' \setminus G \cdot e_0$ is collapsible, the connected component of $X' \setminus G \cdot \mathring{e}_0$ containing η has no valence 1 vertex. Its stabilizer is a cyclic factor H and this connected component is the minimal subtree X'_H. For every $g \in G$, $gX'_H \cap X'_H \neq \varnothing \Rightarrow g \in H$. The subtree X'_H has infinite diameter because it contains η. As X'/G is finite, there exists a vertex with unbounded H-orbit so H is not elliptic. Thus it is not cyclic.

If H is not solvable, then the action of H on X'_H is irreducible: for every segment $I \subset X'_H$, there exists $h \in H$ whose axis contains I.

Let $I \subset \eta$ be a segment containing a $2\,\mathrm{BBT}(\tau)/m$-neighbourhood of μ_j, where $m = \min_{e \notin G \cdot e_0} \frac{\mathrm{len}(\tau(e))}{\mathrm{len}(e)}$. Let $h \in H$ be a loxodromic element whose axis contains I.

The cancellation in $\tau(\mathrm{Axis}_{X'}(h))$ does not reach μ_j so $\mathrm{Axis}_{T_r}(h)$ contains $[\tau(\mu_j)]$.

Finally we must deal with the case where H is isomorphic to $\mathrm{BS}(1, n)$. The subtree X'_H is reduced, so X'_H/H consists in a single edge. If $n = \pm 1$ then X'_H is a line. If $h \in H$ is a loxodromic element then its axis contains η. Moreover since X_H has only valence 2 vertices, they have to belong to different gates so all turns are τ-legal. Therefore $\mathrm{Axis}_{T_r}(h)$ contains $[\tau(\mu_j)]$.

If $|n| > 1$ then X_H is not a line but there is a fixed point ξ in ∂X_H for the action of H. If the line η has both endpoints different from ξ then it might be impossible to find h containing μ_j as a whole. However η contains only one orbit of edge e. Up to reversing the orientation of e we may assume $G_e = G_{t(e)}$, so every turn of the form $\{e, a\bar{e}\}$ with $a \in G_{t(e)}$ is degenerate. Therefore η maps to $\ldots \bar{e}\bar{e}\bar{e} \ldots \bar{e}e \ldots eee \ldots$ in X'_H/H. Since $e_i \ldots e_l$ has length $> 2L$ there exists a subsegment η_0 with length $> L$ of the form $eee \ldots$ or $\bar{e}\bar{e}\bar{e} \ldots$ Once again such a segment is τ-legal otherwise there would only be one gate at the vertices of X_H. There exists a loxodromic element $h \in H$ such that $\mathrm{Axis}_{X'}(h)$

contains η_0. Therefore $\mathrm{Axis}_{X'}(h)$ contains a L-piece of Λ^+ so once again we can apply the discussion above.

We proved that for any $j \in \{1, \dots, k\}$, there exists a translate of e_0 in X' such that the minimal edge subpath μ_j contains a translate of e_0.

If the segments μ_j, $j \in \{1, \dots, k\}$ are disjoint, then we are done. This may fail though. We will see that $\mu_j \cap \mu_{j'} = \varnothing$ if $|j - j'| \geq 2$ and may be a single edge if $|j - j'| = 1$. Thus when counting the translates of e_0 in the $\mu_j, j \in \{1, \dots, k\}$, a translate may be counted more than once, but it can be counted for at most twice (see Figure 3.12). Therefore $\mathrm{Axis}_{X'}(g)/g$ contains at least $k/2$ translates of e_0. Since this lifts to X we get the lemma.

Let us prove the fact about the intersection of the segments μ_j.

First we prove that a $2L$-piece of Λ^+ cannot be contained in the τ-image of a single edge e. By contradiction, suppose otherwise: again we will construct a simple pair of elements containing long pieces of opposite laminations. The edge e must be in the orbit of e_0. As above, one can find τ-legal turns $\{e, e_1\}$ and $\{\bar{e}, e_2\}$ with $e_1, e_2 \notin G \cdot e_0$. Let e' be an edge of $X' \setminus G \cdot e_0$. Since e' is not collapsible, there exists a translate $he' \neq e'$ such that $\{e', he'\}$ (if e' is not a loop in the quotient) or $\{e', h\bar{e}'\}$ (if e' is a loop) is a non-degenerate turn. Define

$$\rho_i := e_i \cdot h_i \bar{e}_i \cdot h_i h_i' e_i \cdot h_i h_i' h_i \bar{e}_1 \dots$$

if e_i is not a loop, where h_i, h_i' are such that $\{e_i, he_i\}$ and $\{\bar{e}_i, h'\bar{e}_i\}$ are non-degenerate. If e_i is a loop define

$$\rho_i := e_i \cdot h_i e_i \cdot h_i^2 e_i \dots$$

Suppose e_i is not a loop. Then e_i cannot be identified with $h_i e_i$ by τ, because the vertex groups at $t(e_i)$ and $h_i t(e_i)$ are not nested: it would imply that T_r has an elliptic element which is not elliptic in X'. There exists a subdivision of X' such that $e_i = a_i \cdot e_i' \cdot b_i$, with $\tau(b_i) = \tau(h_i b_i)$ and $\tau(a_i) = \tau(h_i' a_i)$. Then

$$[\tau(\rho_i)] := \tau(a_i) \cdot \tau(e_i') \cdot h_i \tau(\bar{e}_i') \cdot h_i h_i' \tau(e_i') \dots$$

and since the turns between e and e_i are legal, there is no simplification between $\tau(a_i)$ and $\tau(e)$.

If e_i is a loop, then we already proved in Lemma 3.6.3 that e_i and $h_i \bar{e}_i$ cannot have the same image by τ, or we would obtain new elliptic elements. Once again there exists a

subdivision $e_i = a_i \cdot e'_i \cdot b_i$ with $\tau(\bar{b}_i) = h_i\tau(a_i)$. Thus

$$[\tau(\rho_i)] := \tau(a_i) \cdot \tau(e'_i) \cdot h_i\tau(e'_i) \cdot h_i^2\tau(e'_i) \dots$$

and once again there is no cancellation between $\tau(a_i)$ and $\tau(e)$.

Consequently we can construct a bi-infinite geodesic $\bar{\rho}_1 \cdot e \cdot \rho_2$ such that ρ_1, ρ_2 are rays which cross only one orbit of edges, and $\tau(e) \subset [\tau(\bar{\rho}_1 \cdot e \cdot \rho_2)]$. We proved above that the rays need not be legal, the point is that the cancellation which may occur at turns remains controlled.

Let l be such that the image of any segment of longer than l by τ is longer than $2\,\mathrm{Lip}(\tau)$. Let ρ_1^0, ρ_2^0 be prefixes of the rays longer than l. The path $\bar{\rho}_1^0 \cdot e \cdot \rho_2^0$ can be closed into a loop in the quotient, representing an element $h \in G$ such that $\mathrm{Axis}_{T_r}(h) = [\tau(\mathrm{Axis}_X(h))]$ contains $\tau(e)$, hence a $2L$-piece of Λ^+. The point is that this loop may be constructed such that it crosses only three orbits of edges in X'. Since $b_1(X'/G) \geq 3$, $\mathrm{Axis}_{X'}(h)$ must avoid one orbit of edges.

There exists a candidate u of X' whose axis avoids the same orbit of edge as h, thus $\{u, h\}$ is a simple pair. By assumption $\mathrm{Axis}_{T_r}(u)$ contains an L-piece of Λ^-, which is a contradiction to Proposition 3.3.13: thus $\tau(e)$ cannot contain a whole $2L$-piece of the lamination.

The second point is that the intersection of segments μ_i, μ_j cannot be more than one edge: by minimality of μ_j, the last endpoint of σ_j lies in the image of the last edge of μ_j but not in the image of any other edge. Similarly for $j' > j$, the first point of $\sigma_{j'}$ lies in the image of the first edge of $\mu_{j'}$ and not in any other edge. Thus the intersection $\mu_i \cap \mu_j$ is at most one single edge.

Besides, if $|j' - j| \geq 2$, then σ_{j+1} must be contained between the last point of σ_j and first point of $\sigma_{j'}$. This is not possible if both belong to the same edge, hence the fact. □

Now we have sufficient tools to prove Lemma 3.6.2.

Proof of Lemma 3.6.2. Let X, Y be as in the statement of the lemma. Assume $t_Y < t_X$ in $\mathcal{L}_f$. The other case works similarly by exchanging the roles of ϕ and ϕ^{-1} and will give other constants s, c: we will take the greater constants.

There exists s_1 such that for any $s > s_1$, for any $t \in \mathbb{R}$, for any candidate g of T_t, $\mathrm{Leg}_f(g, T_{t+s}) > \varepsilon_0$. This is a consequence of the following facts: T_t has an f-legal candidate

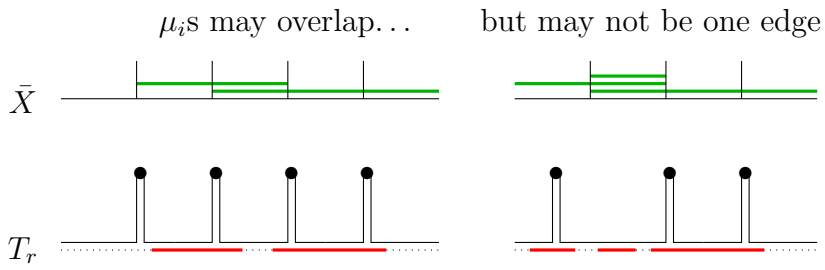

Figure 3.12 – The subsegments of $\mathrm{Axis}_X(g)$ whose image contain a $2L$-piece of Λ^+ may overlap

u (Lemma 3.1.22) so $t_0(u) \leq t$, and there exists $s > 0$ such that for any other candidate v we have $t_0(v) \leq t_0(u) + s$ (Lemma 3.5.4).

Let L be the constant from Proposition 3.3.13. By Lemma 3.4.13 there exists $N_1 \in \mathbb{N}$ such that for every candidate g of X, $\mathrm{LR}(\phi^{N_1}(g), T, \Lambda^+, 2L) > \varepsilon_0/4$.

There exists $s_2 > 0$ such that for any $t \in \mathbb{R}$, the image of any candidate of T_t in T_{t-s_3} contains an L-piece of Λ^-.

Define $d = s_1 + N_1 \log(\lambda) + s_2$. Suppose $t_X - t_Y > d$. Let $r = t_X - s_2$. Let $g \in G$ be a candidate of Y which realizes $\mathrm{Lip}(Y, \pi(X))$. Then the axis of g in T_r contains long leaf segments of Λ^+. Actually $\mathrm{LR}(g, T_r, \Lambda^+, 2L) \geq \varepsilon_0/4$. A given leaf segment of length longer than $2L$ can be at least half covered with disjoint $2L$-pieces of Λ^+. Thus a proportion of at least $\varepsilon_0/8$ of $\mathrm{Axis}_{T_r}(g)$ can be covered by disjoint $2L$-pieces of Λ^+.

Let $k(r)$ be the number of disjoint $2L$-pieces of Λ^+ which tile $\mathrm{Axis}_{T_r}(g)/\langle g \rangle$. We have

$$2Lk(r) > \frac{\varepsilon_0}{8} \|g\|_{T_r}$$

Now $\mathrm{Axis}_{T_r}(g)/\langle g \rangle$ contains $k(r)$ $2L$-pieces of Λ^+. By Lemma 3.6.4 $\mathrm{Axis}_X(g)/\langle g \rangle$ contains at least $k(r)/2$ edges in each orbit of $E(X)$. Since $\mathrm{vol}(X/G) = 1$ we have

$$\mathrm{vol}(\mathrm{Axis}_X(g)/\langle g \rangle) = \|g\|_X \geq k(r)/2.$$

Thus

$$\|g\|_X \geq k(r)/2 \geq \|g\|_{T_r} \frac{\varepsilon_0}{16L}$$

Then

$$\begin{aligned}\mathrm{Lip}(Y,X) &= \frac{\|g\|_X}{\|g\|_Y} \\ &\geq \frac{\|g\|_{T_r}}{\|g\|_Y}\frac{\varepsilon_0}{16L} \\ &\geq \frac{\varepsilon_0}{16L}\mathrm{Lip}(Y,T_r)\end{aligned}$$

By triangular inequality $d_{\mathrm{Lip}}(Y,T_r) \geq d_{\mathrm{Lip}}(Y,\pi(X)) - r$ so

$$d_{\mathrm{Lip}}(Y,X) \geq d_{\mathrm{Lip}}(Y,\pi(X)) - s_2 - \log(\frac{16L}{\varepsilon_0})$$

□

Definition 3.6.5. The ball of outward radius $r > 0$ centered at $Y \in \mathcal{D}$ is

$$B_{\rightarrow}(Y,r) := \{X \in \mathcal{D} / d_{\mathrm{Lip}}(Y,X) < r\}$$

A *closest point projection* to $\mathcal{L}_f$ is a map $p_f : \mathcal{D} \to \mathcal{L}_f$ such that for all $X \in \mathcal{D}$, the distance $d_{\mathrm{Lip}}(X, p_f(X))$ is minimal. The map π_f constructed in Section 3.5 is a closest point projection to $\mathcal{L}_f$.

Now we can state and prove the strong contraction property.

Theorem 3.6.6. *Let ϕ be a fully irreducible automorphism such that ϕ, ϕ^{-1} both admit train track representatives.*

Let $\mathcal{L}_f$ be an axis for ϕ in $\mathcal{D}$ and let p_f be a closest point projection to $\mathcal{L}_f$. Then there exists $D > 0$ such that for any $Y \in \mathcal{D}$ and $r > 0$ such that $B_{\rightarrow}(Y,r) \cap \mathcal{L}_f = \varnothing$

$$\mathrm{diam}(p_f(B_{\rightarrow}(Y,r))) \leq D$$

Proof. Let $Y \in \mathcal{D}$, $r = d_{\mathrm{Lip}}(Y, p_f(Y))$. Let $B := B_{\rightarrow}(Y,r)$. A ball centred at Y intersects the axis if and only if its radius it greater than r. Balls with smaller radius are contained in B so it suffices to bound $\mathrm{diam}(p_f(B))$ independently of Y. Let $X \in B$. Let s, c be the constants from Lemma 3.6.2 and s', c' be the constants from Lemma 3.6.1. Suppose $d_{\mathrm{Lip}}(p_f(Y), p_f(X)) > \max\{s, s'\}$. Then Lemma 3.6.2 yields

$$d_{\mathrm{Lip}}(Y,X) \geq d_{\mathrm{Lip}}(Y, p_f(X)) - c$$

and using Lemma 3.6.1:

$$d_{\mathrm{Lip}}(Y, X) \geq d_{\mathrm{Lip}}(Y, p_f(Y)) + d_{\mathrm{Lip}}(p_f(Y), p_f(X)) - c - c'$$

Since $d_{\mathrm{Lip}}(Y, X) < r = d_{\mathrm{Lip}}(Y, p_f(Y))$ we have

$$d_{\mathrm{Lip}}(p_f(Y), p_f(X)) \leq c + c'$$

Therefore $\mathrm{diam}(p_f(B)) \leq 2\max\{s, s', c + c'\}$. This bound is independent of Y. □

BIBLIOGRAPHY

[Alg11] Yael Algom-Kfir. « Strongly contracting geodesics in outer space ». English. In: *Geom. Topol.* 15.4 (2011), pp. 2181–2233 (cit. on pp. 19, 20, 130–132, 139, 169, 175, 179, 180, 196).

[BS62] Gilbert Baumslag and Donald Solitar. « Some two-generator one-relator non-Hopfian groups ». In: *Bull. Amer. Math. Soc.* 68 (1962), pp. 199–201 (cit. on pp. 7, 25, 73).

[Bee15] Benjamin Beeker. « Multiple conjugacy problem in graphs of free abelian groups ». In: *Groups Geom. Dyn.* 9.1 (2015), pp. 1–27 (cit. on pp. 15, 29, 57).

[BFH97] M. Bestvina, M. Feighn, and M. Handel. « Laminations, trees, and irreducible automorphisms of free groups ». In: *Geom. Funct. Anal.* 7.2 (1997), pp. 215–244 (cit. on pp. 22, 92, 100, 139, 151, 162, 172).

[Bes11] Mladen Bestvina. « A Bers-like proof of the existence of train tracks for free group automorphisms. » English. In: *Fundam. Math.* 214.1 (2011), pp. 1–12 (cit. on p. 143).

[BBF15] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. « Constructing group actions on quasi-trees and applications to mapping class groups ». English. In: *Publ. Math., Inst. Hautes Étud. Sci.* 122 (2015), pp. 1–64 (cit. on pp. 20, 131).

[BF91] Mladen Bestvina and Mark Feighn. « Bounding the complexity of simplicial group actions on trees ». In: *Invent. Math.* 103.3 (1991), pp. 449–469 (cit. on pp. 34, 35, 50, 157).

[BF14] Mladen Bestvina and Mark Feighn. « Hyperbolicity of the complex of free factors ». English. In: *Adv. Math.* 256 (2014), pp. 104–155 (cit. on p. 10).

[BH92] Mladen Bestvina and Michael Handel. « Train tracks and automorphisms of free groups ». In: *Ann. of Math. (2)* 135.1 (1992), pp. 1–51 (cit. on pp. 17, 18, 74, 75, 86, 100, 131, 136).

[Bou16] Margot Bouette. « On the growth of the automorphisms of Baumslag-Solitar groups ». Theses. Université Rennes 1, Dec. 2016 (cit. on pp. 14, 19, 26, 75, 89, 131).

[Cla09] Matt Clay. « Deformation spaces of G-trees and automorphisms of Baumslag-Solitar groups ». In: *Groups Geom. Dyn.* 3.1 (2009), pp. 39–69 (cit. on pp. 9, 25).

[CF08a] Matt Clay and Max Forester. « On the isomorphism problem for generalized Baumslag-Solitar groups ». In: *Algebr. Geom. Topol.* 8.4 (2008), pp. 2289–2322 (cit. on pp. 9, 25).

[CF08b] Matt Clay and Max Forester. « On the isomorphism problem for generalized Baumslag-Solitar groups ». English. In: *Algebr. Geom. Topol.* 8.4 (2008), pp. 2289–2322.

[Col78] Donald J. Collins. « The automorphism towers of some one-relator groups ». English. In: *Proc. Lond. Math. Soc. (3)* 36 (1978), pp. 480–493 (cit. on p. 9).

[CL83] Donald J. Collins and Frank Levin. « Automorphisms and Hopficity of certain Baumslag-Solitar groups ». In: *Arch. Math. (Basel)* 40.5 (1983), pp. 385–400 (cit. on pp. 7, 9, 25).

[Coo87] Daryl Cooper. « Automorphisms of free groups have finitely generated fixed point sets ». English. In: *J. Algebra* 111 (1987), pp. 453–456 (cit. on pp. 85, 140).

[CL15] Thierry Coulbois and Martin Lustig. « Long turns, INP's and indices for free group automorphisms ». In: *Illinois J. Math.* 59.4 (2015), pp. 1087–1109 (cit. on p. 104).

[CM87] Marc Culler and John W. Morgan. « Group actions on $\mathbb{R}$-trees ». English. In: *Proc. Lond. Math. Soc. (3)* 55 (1987), pp. 571–604 (cit. on p. 125).

[CV86] Marc Culler and Karen Vogtmann. « Moduli of graphs and automorphisms of free groups ». English. In: *Invent. Math.* 84 (1986), pp. 91–119.

[For02] Max Forester. « Deformation and rigidity of simplicial group actions on trees. » English. In: *Geom. Topol.* 6 (2002), pp. 219–267 (cit. on pp. 11, 139).

[For06] Max Forester. « Splittings of generalized Baumslag-Solitar groups ». In: *Geom. Dedicata* 121 (2006), pp. 43–59 (cit. on pp. 11, 14, 25, 26, 32, 73, 75, 80, 131).

[FM15] Stefano Francaviglia and Armando Martino. « Stretching factors, metrics and train tracks for free products ». English. In: *Ill. J. Math.* 59.4 (2015), pp. 859–899 (cit. on pp. 17, 143, 145, 146).

[Gab+98] Damien Gaboriau et al. « An index for counting fixed points of automorphisms of free groups ». English. In: *Duke Math. J.* 93.3 (1998), pp. 425–452 (cit. on p. 139).

[Ger87] S. M. Gersten. « Fixed points of automorphisms of free groups ». English. In: *Adv. Math.* 64 (1987), pp. 51–85 (cit. on p. 76).

[Gil+00] N. D. Gilbert et al. « Tree actions of automorphism groups ». English. In: *J. Group Theory* 3.2 (2000), pp. 213–223 (cit. on p. 9).

[Gui04] Vincent Guirardel. « Limit groups and groups acting freely on $\mathbb{R}^n$-trees. » English. In: *Geom. Topol.* 8 (2004), pp. 1427–1470 (cit. on p. 98).

[GH19] Vincent Guirardel and Camille Horbez. « Algebraic laminations for free products and arational trees ». In: *Algebr. Geom. Topol.* 19.5 (2019), pp. 2283–2400 (cit. on pp. 27, 28, 43).

[GL07] Vincent Guirardel and Gilbert Levitt. « Deformation spaces of trees ». In: *Groups Geom. Dyn.* 1.2 (2007), pp. 135–181 (cit. on pp. 12, 31–33, 66, 78, 85, 121, 134, 135, 147).

[GL17] Vincent Guirardel and Gilbert Levitt. « JSJ decompositions of groups ». In: *Astérisque* 395 (2017), pp. vii+165 (cit. on p. 51).

[HV98] Allen Hatcher and Karen Vogtmann. « The complex of free factors of a free group ». English. In: *Q. J. Math., Oxf. II. Ser.* 49.196 (1998), pp. 459–468 (cit. on p. 10).

[Kap14] Ilya Kapovich. « Algorithmic detectability of iwip automorphisms ». In: *Bull. Lond. Math. Soc.* 46.2 (2014), pp. 279–290 (cit. on pp. 17, 75, 78, 95, 196).

[Kap19] Ilya Kapovich. « Detecting fully irreducible automorphisms: a polynomial time algorithm ». In: *Exp. Math.* 28.1 (2019). With an appendix by Mark C. Bell, pp. 24–38 (cit. on p. 75).

[Lev07] Gilbert Levitt. « On the automorphism group of generalized Baumslag-Solitar groups ». In: *Geom. Topol.* 11 (2007), pp. 473–515 (cit. on pp. 9, 25, 31, 75, 80, 89, 119).

[LS01] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory*. Classics in Mathematics. Reprint of the 1977 edition. Springer-Verlag, Berlin, 2001, pp. xiv+339 (cit. on p. 28).

[MM99] Howard A. Masur and Yair N. Minsky. « Geometry of the complex of curves. I: Hyperbolicity ». English. In: *Invent. Math.* 138.1 (1999), pp. 103–149 (cit. on p. 10).

[Mei15] Sebastian Meinert. « The Lipschitz metric on deformation spaces of G-trees ». In: *Algebr. Geom. Topol.* 15.2 (2015), pp. 987–1029 (cit. on pp. 17, 26, 75, 86, 131, 142, 143, 145, 147).

[Mes72] Stephen Meskin. « Nonresidually finite one-relator groups ». In: *Trans. Amer. Math. Soc.* 164 (1972), pp. 105–114.

[Min96] Yair N. Minsky. « Quasi-projections in Teichmüller space ». English. In: *J. Reine Angew. Math.* 473 (1996), pp. 121–136.

[MS84] John W. Morgan and Peter B. Shalen. « Valuations, trees, and degenerations of hyperbolic structures. I ». English. In: *Ann. Math. (2)* 120 (1984), pp. 401–476 (cit. on p. 12).

[Pau89] Frédéric Paulin. « The Gromov topology on $\mathbb{R}$-trees ». English. In: *Topology Appl.* 32.3 (1989), pp. 197–221 (cit. on p. 12).

[Sen81] E. Seneta. *Non-negative matrices and Markov chains. 2nd ed.* English. Springer, New York, NY, 1981 (cit. on pp. 88, 137).

[Ser77] Jean-Pierre Serre. *Arbres, amalgames,* SL_2. Avec un sommaire anglais, Rédigé avec la collaboration de Hyman Bass, Astérisque, No. 46. Société Mathématique de France, Paris, 1977, 189 pp. (1 plate) (cit. on pp. 7, 10, 30, 54, 56, 57, 81, 83, 116, 132, 133).

[Sou01] E. Souche. « Quasi-isométries et quasi-plans dans l'étude des groupes discrets ». In: 2001.

[Sta91] John R. Stallings. « Foldings of G-trees ». In: *Arboreal group theory (Berkeley, CA, 1988)*. Vol. 19. Math. Sci. Res. Inst. Publ. Springer, New York, 1991, pp. 355–368 (cit. on p. 34).

[Whi36] J. H. C. Whitehead. « On equivalent sets of elements in a free group ». In: *Ann. of Math. (2)* 37.4 (1936), pp. 782–800 (cit. on p. 27).

UNIVERSITÉ DE RENNES 1

Titre : Facteurs cycliques, automorphismes irréductibles et axes contractants pour les groupes de Baumslag-Solitar généralisés

Mot clés : automorphismes de groupes de Baumslag-Solitar, espaces de déformation, représentants train track, facteurs invariants, construction d'espaces hyperboliques

Résumé : Un *groupe de Baumslag-Solitar généralisé* est le groupe fondamental d'un graphe de groupes cycliques infinis. Cette thèse porte sur l'étude du groupe d'automorphismes extérieurs d'un tel groupe G. Par analogie avec l'outre-espace associé au groupe d'automorphismes extérieurs $Out(F_N)$, on exploite l'action de $Out(G)$ sur un *espace de déformation*, qui est un espace dont les points sont des actions de G sur des arbres. Nous définissons les *facteurs cycliques* comme les facteurs qui apparaissent dans les scindements cycliques de G, et donnons un algorithme qui décide si un élément donné du groupe G appartient à un tel facteur. Les *automorphismes complètement irréductibles* sont des automorphismes dont aucune puissance ne fixe la classe de conjugaison d'un tel facteur. Nous donnons un algorithme analogue à [Kap14] qui permet de décider si un automorphisme est irréductible lorsque celui-ci a un train track et aucun élément pseudo-périodique. Enfin, inspirée par [Alg11], nous montrons que les projections au point le plus proche sur les axes de translation dans l'espace de déformation des automorphismes irréductibles admettant des train track sont fortement contractantes.

Title: Cyclic factors, irreducible automorphisms and contracting axes for generalized Baumslag-Solitar groups

Keywords: Automorphisms of Baumslag-Solitar groups, deformation spaces, train track representatives, invariant factors, construction of hyperbolic spaces

Abstract: A *generalized Baumslag-Solitar group* is the fundamental group of a graph of infinite cyclic groups. This thesis focuses on the outer automorphism group of such a group G. In analogy with the Outer Space associated to the outer automorphism group $Out(F_N)$, we study the action of $Out(G)$ on a *deformation space* which is a space whose points are some actions of G on trees. We define *cyclic factors* as subgroups which occur in splittings over cyclic groups, and give an algorithm deciding if a given element of G belongs to such a factor. *Fully irreducible automorphisms* are automorphisms having no periodic conjugacy class of cyclic factor. We give an algorithm which decides whether an automorphism is irreducible, given that it has a train track representative and no periodic element; it is analogous to [Kap14]. Finally, inspired by [Alg11], we prove that closest point projections to the axis of an irreducible element with a train track representative are strongly contracting.