

Université de Caen/Basse-Normandie
U.F.R. de Sciences
École doctorale SIMEM



T H È S E

présentée par

Chloé PERIN

et soutenue le vendredi 31 octobre 2008

en vue de l'obtention du

DOCTORAT de l'UNIVERSITÉ de CAEN

Spécialité : mathématiques et leurs interactions
(Arrêté du 07 août 2006)

Plongements élémentaires dans un groupe hyperbolique sans torsion

MEMBRES du JURY

- M. Tuna ALTINEL, maître de conférences (H.D.R) à l'Université de Lyon 1
- M. Thomas DELZANT, professeur à l'Université de Strasbourg 1 (*rapporteur*)
- M. Vincent GUIRADEL, maître de conférences à l'Université de Toulouse 3
- M. Gilbert LEVITT, professeur à l'Université de Caen (*directeur*)
- M. Zlil SELA, professeur à l'Université Hébraïque de Jérusalem

Remerciements

Je souhaite remercier en premier lieu les deux personnes qui m'ont guidée au cours de ces trois ans : Gilbert Levitt et Zlil Sela. Tous deux ont fait preuve d'une grande disponibilité et d'une infinie patience face à mes questions. J'ai pu apprécier tout particulièrement la facilité de Gilbert pour improviser une explication au tableau, déroulant devant mes petits yeux ébahis tout un pan de mathématiques dont je n'avais pas soupçonné l'existence. Son attention au détail, et sa persévérance pour résoudre avec moi des problèmes techniques pas toujours passionnants m'ont été d'une aide précieuse. Quant à Zlil, il a su me faire découvrir son univers mathématique et me révéler peu à peu les divers aspects de ce domaine passionnant. Sa gentillesse et sa simplicité ont grandement contribué à me mettre à l'aise dès mon premier séjour à Jérusalem.

J'exprime également toute ma reconnaissance aux deux rapporteurs, Thomas Delzant et Panos Papasoglu, qui ont énormément contribué à l'amélioration de ce texte par leurs remarques et leurs conseils. Je remercie tout particulièrement Thomas pour les quelques jours qu'il m'a consacrés à Strasbourg, et pour un dîner de tartes flambées en famille très sympathique. Quant à Panos, qui ne peut malheureusement pas être ici aujourd'hui, je le remercie d'avoir accepté cette lourde tâche, prouvant ainsi qu'il ne m'a pas gardé rancune pour l'avoir battu au pendu à Berkeley.

Je remercie également les autres membres du jury : Tuna Altinel qui apporte au jury le point de vue indispensable d'un logicien, et Vincent Guirardel, qui a été de bon conseil à plusieurs reprises ces trois dernières années.

Je profite de cette occasion pour remercier tous les professeurs qui m'ont inspirée, du lycée à la thèse et grâce à qui (à cause de qui ?) j'ai choisi de faire de la recherche en mathématiques : Mme Nyers, Keith Carne et Martin Hyland (MHLF), et enfin Frédéric Paulin et Pierre Pansu. Je remercie tout particulièrement ces deux derniers pour m'avoir orientée lors de mon choix de sujet de thèse, et pour m'avoir mise en contact avec Gilbert et Zlil.

J'ai aussi une pensée particulière pour tous les gens avec qui j'ai découvert le plaisir de faire des maths à plusieurs : dans un jardin, dans une rame de RER, dans un café, à table, à 4h du matin la veille d'une supervision, sur un coin de nappe ou un ticket de caisse, avec les mains ou au téléphone. David et Adam à King's ; Katharina pendant la rédaction de notre mémoire de maîtrise commun (sur la preuve de la conjecture de géométrisation... en dimension 1 !); Erwan et Dimitri au DEA de logique ; Benoît, Nico, Philippe et sa quête éternelle pour la démonstration parfaite pendant la préparation à l'agrégation : tous m'ont montré, preuve à l'appui, que les maths sont **aussi** une activité sociale ! Nicolas et Xavier seraient aussi dans cette liste, s'ils s'étaient orientés vers un domaine raisonnable des mathématiques. Faute de quoi, nous avons dû trouver d'autres sujets de conversation, mais finalement je ne suis pas certaine d'avoir perdu au change.

Pendant ces trois années, j'ai eu le plaisir de côtoyer plusieurs générations de thésards de Caen. Je remercie particulièrement les vieux sages qui m'ont accueillie dans la bande et montré les ficelles du métier de thésard (à quelle heure aller au RU ? à qui s'adresser en cas de problème informatique ? où trouver un fichier `.tex` à repiquer pour faire son projet professionnel ?). Mais je n'oublie pas mes contemporains, ainsi que la jeune génération.

Et puis bien sûr, il y a Marc. J'envahis son appartement trois mois par an, je casse ses ustensiles

de cuisine, je l'empêche de travailler et je monopolise son téléphone, tout ça pour repartir pendant six mois sans donner de nouvelles et recommencer l'année suivante. Et il supporte ça avec le sourire. Mes parents aussi, d'accord. Mais Marc, lui, résout aussi mes problèmes de \LaTeX . Il a donc été d'une aide irremplaçable, mais surtout, sans nos discussions à bâtons rompus, sans nos humours accordés et nos horaires décalés, ces trois années n'auraient pas eu la même saveur.

Je veux aussi remercier ma famille, à qui je dois tout : mes parents, dont la tolérance face à mon choix de carrière les honore, et Arthur, Dorothée et Marguerite, qui m'acceptent telle que je suis même devant leurs amis branchés. Leurs efforts répétés pour retenir le sujet de ma thèse mérite une mention particulière. Mais surtout, leurs stratégies variées de soutien moral, qu'elles soient postales, téléphoniques, théâtrales ou même chorégraphiques m'ont fait passer plus d'un cap difficile avec le sourire (voire le fou rire).

Enfin, pour son soutien inconditionnel et sa partialité décidée à mon égard, je remercie Eran. Je suis heureuse d'être non seulement celle avec qui il a choisi de partager sa vie, mais aussi sa revanche ultime sur le prof de maths qui l'a renvoyé de son cours au lycée.

Acknowledgements

First of all, I would like to thank warmly the two people who guided and advised me during these three years: Gilbert Levitt and Zlil Sela. Both have been very present, and have shown great patience in answering my many questions. The ability of Gilbert to improvise explanations and surveys of whole theories on the blackboard has often left me filled with wonder. His attention to detail, and his tenacity in helping me solve many technical problems have been of great help. As for Zlil, he introduced me to his mathematical world and patiently revealed for me some of the wonders of this fascinating domain. His kindness and simplicity have greatly contributed to make me feel at home from my first day at the Hebrew University.

I would like to express my gratitude to the referees, Thomas Delzant and Panos Papasoglu, who contributed greatly to the improvement of this text by their comments. I particularly thank Thomas for the few days he spent working with me in Strasbourg, and for a very nice family dinner. As for Panos, who can unfortunately not be here today, I am grateful to him for taking on this arduous task, without bearing me grudge for defeating him at hangman in Berkeley.

I also thank the other members of the jury: Tuna Altinel who adds to the jury the much needed perspective of a logician, and Vincent Guirardel, whose helpful advice was regularly sought and found over the past three years.

This is a good occasion to thank some of the professors who inspired me, from high school up until today, and thanks to whom (because of whom?) I chose this career path: Mme Nyers, Keith Carne and Martin Hyland (MHLF), and finally Frédéric Paulin and Pierre Pansu. I particularly thank them both for guiding me in my choice of a PhD subject, and for putting me in touch with Gilbert and Zlil.

I would also like to mention some of the people with whom I shared the pleasure of doing maths: in a park, in a train or in a pub, at 4am the night before a supervision, on a napkin, by hand-waving or on the phone. David and Adam at King's; Katharina during our common Master thesis (on the proof of the geometrisation conjecture... in dimension 1!); Erwan and Dimitri in the Logic Master classes; and the year of the agrégation, Benoît, Nico, and Philippe and his eternal quest for the perfect proof: all of them showed me that mathematics are **also** a social activity! Nicolas and Xavier would have made it into the above list, had they opted for a sensible area of mathematics. We have thus been forced to come up with other discussion topics, but on the whole I am not sure I regret it.

During these three years, it has been my pleasure to mingle with several generation of PhD students of the University of Caen. I especially thank my elders who welcomed me and taught me the basic survival techniques of a PhD student (what time is best to go to the cafeteria? who should I turn to regarding computing issues?). But I do not forget my contemporaries, as well as the younger generations.

And of course, there is Marc. I take over his flat three months every year, I destroy his cutlery, I prevent him from working and I block his phone line, after which I disappear for six month without a word just to turn up again the following year. And Marc graciously bears it. Right, so do my parents. But Marc also solves my L^AT_EX problems. His help has thus been

irreplaceable, but more importantly, without our endless discussions, without our tuned sense of humour and our desynchronised schedules, these three years would simply not have been the same.

I also wish to thank my family, to whom I owe everything: my parents for being remarkably tolerant towards my disreputable career choice, and Arthur, Dorothée and Marguerite for accepting me just as I am even in front of their cool friends. Their repeated efforts for remembering the topic of my PhD deserves at least a "Well Done" sticker. Most importantly, their inventive moral support initiatives, be they postal, telephonic, theatrical or even choreographic have helped me go through the inevitable rough patches with a smile on my face, if not crying with laughter.

Last but not least, I will never thank Eran enough for his unconditional support and his categorical partiality in all things concerning me. I am happy to be both the person he chose to share his life with, and his sweet revenge on the math teacher who kicked him out of his class.

Table des matières

1	Introduction	1
1.1	Plongements élémentaires	2
1.2	Tours hyperboliques	5
1.3	Structure de la démonstration	6
1.3.1	Construction du morphisme f	7
1.3.2	Démonstration de la proposition C	8
1.4	Contenu de la thèse	9
2	Introduction in English	11
2.1	Elementary embeddings	12
2.2	Hyperbolic towers	14
2.3	Structure of the proof	16
2.3.1	Construction of the morphism f	16
2.3.2	Proof of Proposition C	18
2.4	Content of the thesis	18
3	Préliminaires	21
3.1	Actions sur des arbres simpliciaux et graphes de groupes	21
3.2	Espaces métriques hyperboliques	23
3.3	Limites d'espaces métriques	27
3.3.1	La topologie de Gromov-Hausdorff	27
3.3.2	Ultraproduits et limites	28
3.3.3	Limites de G -espaces pointés hyperboliques	30
4	Argument du raccourcissement	31
4.1	Limite d'une suite de morphismes vers un groupe hyperbolique	32
4.2	Raccourcissement des morphismes - le cas classique	35
4.2.1	Groupe modulaire	35
4.2.2	Raccourcissement des actions	36
4.2.3	Raccourcissement des morphismes	38
4.3	Raccourcissement des morphismes - le cas relatif	40
4.3.1	Raccourcissement des actions	40
4.3.2	Raccourcissement des morphismes	41
4.4	La propriété de Co-Hopf relative pour les groupes hyperboliques sans torsion	43
5	Démonstration du théorème de raccourcissement des actions	45
5.1	Exemples d'actions sur des arbres réels	45
5.2	Graphes d'actions	47
5.3	Décomposition de Rips	48

5.4	Cas surface	49
5.5	Cas axial	53
5.6	Cas simplicial	54
5.7	Démonstration du théorème de raccourcissement	61
6	Ensemble de factorisation	63
6.1	Cas des groupes libres	63
6.2	Cas des groupes hyperboliques sans torsion	65
6.3	Cas relatif	67
7	Plongements élémentaires dans un groupe hyperbolique	69
7.1	Tours hyperboliques et énoncé du résultat principal	69
7.2	Décompositions de type JSJ et prérétractions	70
7.2.1	Décompositions de type JSJ	70
7.2.2	Décompositions JSJ	70
7.2.3	Prérétractions	72
7.3	Construction d'une prérétraction à l'aide de la logique du premier-ordre	72
7.4	Démonstration du résultat principal	76
7.5	Le cas particulier des groupes libres	78
8	Une propriété des décompositions de type JSJ	79
9	Applications non pinçantes et propriété de l'indice fini	83
9.1	Surfaces à bord	83
9.1.1	Surfaces agissant sur des arbres simpliciaux	83
9.1.2	Applications non pinçantes et propriété de l'indice fini	84
9.1.3	Complexités	85
9.2	Graphes de groupes à surfaces	86
9.2.1	Raffinements elliptiques de graphes de groupes à surfaces	86
9.2.2	Applications non pinçantes sur des graphes de groupes à surfaces	86
9.2.3	Complexité des surfaces d'un graphe de groupes	87
9.3	Propriété de l'indice fini pour un produit libre	88
10	Des prérétractions aux étages de tour hyperbolique	91
10.1	Pincements	91
10.2	Non-abélianité des surfaces	94
10.3	Prérétractions maximales	96
10.4	Démonstration de la proposition 7.15	99
10.5	Démonstration de la proposition 7.16	101

Contents

1	Introduction	1
1.1	Plongements élémentaires	2
1.2	Tours hyperboliques	5
1.3	Structure de la preuve du théorème B	6
1.3.1	Construction du morphisme f	7
1.3.2	Preuve de la proposition C	8
1.4	Contenu de la thèse	9
2	Introduction in English	11
2.1	Elementary embeddings	12
2.2	Hyperbolic towers	14
2.3	Structure of the proof of theorem B	16
2.3.1	Construction of the morphism f	17
2.3.2	Proof of Proposition C	18
2.4	Content of the thesis	18
3	Basic notions	21
3.1	Actions on simplicial trees and graphs of groups	21
3.2	Hyperbolic metric spaces	23
3.3	Limits of metric spaces	27
3.3.1	The Gromov-Hausdorff topology	27
3.3.2	Ultraproducts and limit of sequences	28
3.3.3	Limits of pointed hyperbolic G -spaces	30
4	Shortening argument	31
4.1	Limit of morphisms to a hyperbolic group	32
4.2	Shortening morphisms in the classical case	35
4.2.1	Modular group	35
4.2.2	Action shortening result	36
4.2.3	Morphism shortening result	38
4.3	Shortening morphisms in the relative case	40
4.3.1	Action shortening result	40
4.3.2	Morphism shortening result	41
4.4	The relative Co-Hopf property for hyperbolic groups	43
5	Proof of the action shortening Theorem	45
5.1	Some examples of actions on real trees	45
5.2	Graphs of actions	47
5.3	Rips decomposition	48

5.4	Surface case	49
5.5	Axial Case	53
5.6	Simplicial Case	54
5.7	Proof of the shortening Theorem	61
6	Factor sets	63
6.1	Case of free groups	63
6.2	Case of torsion-free hyperbolic groups	65
6.3	Relative results	67
7	Elementary embeddings in a hyperbolic group	69
7.1	Hyperbolic towers and statement of the main result	69
7.2	JSJ-like decompositions and preretractions	70
7.2.1	JSJ-like decomposition	70
7.2.2	JSJ decompositions	70
7.2.3	Preretractions	72
7.3	Using first order to build preretractions	73
7.4	Proof of the main result	77
7.5	The special case of free groups	78
8	A property of JSJ-like decompositions	79
9	Non-pinching maps and the finite index property	83
9.1	Surfaces with boundary	83
9.1.1	Surfaces acting on simplicial trees	83
9.1.2	Non-pinching maps and the finite index property	84
9.1.3	Complexities	85
9.2	Graphs of groups with surfaces	86
9.2.1	Elliptic refinements of graphs of groups with surfaces	86
9.2.2	Non-pinching maps on graphs of groups with surfaces	86
9.2.3	Surface complexity of graphs of groups	87
9.3	Finite index property for free products	88
10	From preretractions to hyperbolic floors	91
10.1	Pinching a set of curves	91
10.2	Non-abelianity of surfaces	93
10.3	Maximal preretractions	95
10.4	Proof of Proposition 7.15	99
10.5	Proof of Proposition 7.16	100

Chapitre 1

Introduction

La théorie géométrique des groupes, développée au cours des vingt dernières années, a permis grâce à l'application de concepts de géométrie classique comme la courbure de résoudre de nombreux problèmes de théorie des groupes jusqu'alors abordés d'un point de vue combinatoire ou purement algébrique. La représentation d'un groupe par un espace métrique permet l'apparition de notions telles que la quasi-isométrie entre deux groupes, et de nouvelles classes de groupes comme celle des groupes hyperboliques ou encore des groupes $CAT(0)$. Les propriétés géométriques de l'espace associé à un groupe, au-delà de leur intérêt intrinsèque, s'avèrent être profondément liées aux propriétés algébriques classiques du groupe en question.

Plus récemment, cette approche géométrique s'est révélée particulièrement fructueuse pour l'étude de questions empruntées à la logique, notamment à la théorie des modèles. Cette nouvelle impulsion trouve ses fondements principalement dans les divers travaux effectués sur le problème de Tarski, notamment par Zlil Sela dans [Sel01] ainsi que dans ses autres articles de la même série (voir aussi les travaux de Kharlampovich et Myasnikov [KM06]). L'approche de Sela utilise de manière extensive la théorie de Bass-Serre sur les actions de groupes sur les arbres simpliciaux, la théorie des décompositions JSJ, qui permet de décrire toutes les actions d'un groupe donné sur un arbre simplicial, ainsi que la théorie de Rips, qui analyse des actions de groupes sur un arbre réel. Les résultats obtenus révèlent une corrélation significative entre les propriétés géométriques d'un groupe et sa théorie du premier ordre : Sela montre par exemple qu'un groupe de type fini élémentairement équivalent à un groupe hyperbolique sans torsion est lui-même hyperbolique sans torsion (voir [Sel]).

Ce point de vue géométrique a donc permis de résoudre plusieurs problèmes difficiles de théorie des modèles. Les outils développés pour comprendre ces questions de logique ont considérablement enrichi en retour la théorie géométrique des groupes. Par ailleurs, d'autres problèmes qui peuvent sembler extérieurs aux deux domaines ont pu être résolus grâce aux techniques développées : on peut penser à l'étude des groupes ω -résiduellement libres, ou encore à la résolution d'équations sur les groupes libres. Par exemple, Sela obtient dans [Sel01] une nouvelle preuve de l'existence d'un diagramme de Makanin-Razborov (le résultat original est obtenu par Razborov dans [Raz85] et généralisé par Kharlampovich et Myasnikov dans [KM98]), diagramme qui permet de classer les homomorphismes d'un groupe de type fini dans un groupe libre. La nature géométrique de la preuve lui permet ensuite dans [Sel] une généralisation du résultat aux morphismes d'un groupe de type fini dans un groupe hyperbolique sans torsion. Ceci a encore été généralisé par Groves aux groupes hyperboliques relativement à une collection de groupes abéliens libres dans [Gro05].

Une notion de base de la théorie des modèles est celle de plongement élémentaire, qui décrit comment une structure se plonge dans une autre, de manière indiscernable pour la logique du

premier ordre. Dans cette thèse, on s'intéresse aux groupes élémentairement plongés dans un groupe hyperbolique sans torsion.

1.1 Plongements élémentaires

Pour une introduction rapide à la théorie des modèles, on pourra consulter [Cha]. On appelle **langage des groupes** l'ensemble de symboles

$$\mathcal{L} = \{=, (,), \neg, \vee, \wedge, \forall, \exists, 1, *, ^{-1}\} \cup V,$$

où V est un ensemble de variables infini dénombrable. On rappelle que le symbole \vee représente la disjonction (« ou »), le symbole \wedge la conjonction (« et »), et le symbole \neg la négation (« non »). Le symbole 1 représentera l'élément unité du groupe, $*$ est la loi de composition (mais on s'autorisera à représenter le produit par la concaténation), et $^{-1}$ permet d'exprimer l'inversion. Une **formule du premier ordre** (ou **formule élémentaire**) dans le langage \mathcal{L} est une suite finie d'éléments de \mathcal{L} qui constituent une formule mathématique « grammaticalement correcte ». Par la suite, nous utiliserons fréquemment les symboles mathématiques usuels pour représenter une suite de symboles de \mathcal{L} , comme le symbole \rightarrow , où $A \rightarrow B$ représente $B \vee \neg A$.

On dit qu'une variable x apparaissant dans une formule du premier ordre est **libre** si elle n'est précédée ni de $\forall x$ ni de $\exists x$. Une formule du premier ordre ϕ est dite **close** (on dit aussi que ϕ est un **énoncé**) si aucune des variables apparaissant dans la formule n'est libre. Un groupe G **satisfait** une formule close ϕ dans le langage \mathcal{L} si l'interprétation de la formule est vraie dans G . On note alors $G \models \phi$.

Exemple 1.1 : Si ϕ est la formule $\forall x \forall y x * y * x^{-1} * y^{-1} = 1$, un groupe G satisfait ϕ si et seulement s'il est abélien.

Définition 1.2 : (théorie élémentaire) *Soit G un groupe. La théorie élémentaire de G dans \mathcal{L} est l'ensemble des formules closes du premier ordre dans \mathcal{L} que G satisfait.*

Il est important de remarquer que l'on ne peut quantifier que sur les éléments du groupe, et pas sur les sous-parties, ni sur les entiers naturels.

Exemple 1.3 : Pour exprimer qu'un groupe est sans torsion, on pourrait écrire la formule suivante

$$\forall x (x \neq 1) \rightarrow \bigwedge_{n=1}^{\infty} (x^n \neq 1)$$

Cependant, celle-ci **n'est pas** une formule du premier ordre, puisque si on devait l'écrire en utilisant uniquement les symboles de \mathcal{L} (sans raccourcis), on aurait une formule infinie. De la même manière, la formule

$$\forall x (x \neq 1) \rightarrow \forall n \in \mathbb{N} (x^n \neq 1)$$

n'est pas une formule du premier ordre puisqu'on quantifie sur les entiers.

Définition 1.4 : (élémentairement équivalent) *Deux groupes G et G' sont élémentairement équivalents s'ils ont la même théorie élémentaire dans le langage des groupes. On note alors $G \equiv G'$.*

Exemple 1.5 : Soient G et G' deux groupes élémentairement équivalents.

- Si G est abélien, G' l'est aussi.
- Si G est fini, G' aussi, et ils ont le même cardinal. En fait, ils sont isomorphes : la table de multiplication de G peut être exprimée par une formule du premier ordre, que G' satisfait.

- Si G est sans torsion, G' aussi. Bien que « être sans torsion » ne peut pas être exprimé par une formule du premier ordre, mais peut être exprimé par la famille infinie d'énoncés

$$\{\forall x (x \neq 1 \rightarrow x^n \neq 1)\}_{n \in \mathbb{N} - \{0\}}.$$

Si G est sans torsion, il satisfait chacune de ces formules, donc G' aussi.

Exemple 1.6: Les groupes \mathbb{Z} et \mathbb{Z}^2 ne sont pas élémentairement équivalents. En effet, \mathbb{Z} satisfait

$$\exists x \forall y \exists z (y = z^2) \vee (y = z^2 x),$$

qui exprime que dans \mathbb{Z} , un élément est soit pair soit impair. Clairement, \mathbb{Z}^2 ne satisfait pas cet énoncé. On peut en fait montrer de cette façon que $\mathbb{Z}^k \equiv \mathbb{Z}^l$ si et seulement si $k = l$.

On peut maintenant énoncer le problème suivant :

Question 1: *Supposons $1 < m < n$. Les groupes libres de rang m et n sont-ils élémentairement équivalents ?*

Ce problème, posé par le logicien Alfred Tarski dans les années 40, est connu sous le nom de problème de Tarski. Sela a répondu de manière positive à cette question dans [Sel06]. Les travaux de Kharlampovich et Myasnikov proposent une autre approche de ce problème (voir [KM06]). Sela donne également une caractérisation de tous les groupes de type fini élémentairement équivalents aux groupes libres (voir Théorème 1.15). Le lien avec la géométrie se manifeste de manière frappante dans le résultat suivant, qui est un corollaire de cette caractérisation :

Théorème 1.7 : *Le groupe fondamental d'une surface fermée de caractéristique d'Euler au plus -2 est élémentairement équivalent à un groupe libre de type fini non abélien.*

On peut s'intéresser dans le cadre des groupes libres à d'autres notions classiques de la théorie des modèles, comme celle de sous-structure élémentaire.

Définition 1.8 : (plongement élémentaire) *Soit G un groupe et soit H un sous-groupe de G . On note \mathcal{L}_H le langage des groupes \mathcal{L} auquel on ajoute pour tout élément h de H une nouvelle constante $[h]$. On dit que le plongement $H \subseteq G$ est élémentaire, ou encore que H est un sous-groupe élémentaire de G si pour tout énoncé du premier ordre ϕ dans le langage \mathcal{L}_H , le sous-groupe H satisfait ϕ si et seulement si G satisfait ϕ . On note alors $H \preceq G$.*

Remarquons que cette définition équivaut à dire que H et G sont élémentairement équivalents dans le langage \mathcal{L}_H , et donc entraîne l'équivalence élémentaire classique (dans le langage \mathcal{L}).

Exemple 1.9: Soit G un groupe et soit H un sous-groupe de G . Soit h un élément de H . Considérons l'énoncé

$$\phi_h : \forall x ([h], x) = 1.$$

C'est un énoncé du premier ordre dans le langage \mathcal{L}_H . Le groupe H (respectivement G) satisfait ϕ_h si et seulement si h appartient au centre $Z(H)$ de H (respectivement $Z(G)$ de G). En particulier, si $H \preceq G$, on voit que $h \in Z(H)$ si et seulement si $h \in Z(G)$ et on en déduit $Z(H) = H \cap Z(G)$.

Lorsque l'on s'intéresse à la théorie du premier ordre des groupes libres, la question suivante est naturelle :

Question 2: *Décrire les plongements élémentaires dans un groupe libre.*

Pour montrer l'équivalence élémentaire des groupes libres de type fini, Sela montre en fait le

Théorème 1.10 : [Sel06, Theorem 4] *Soit $i : \mathbb{F}_k \rightarrow \mathbb{F}_n$ le plongement canonique d'un groupe libre à k générateurs dans un groupe libre à n générateurs pour $2 \leq k \leq n$. Alors i est un plongement élémentaire.*

Il est donc naturel de se demander si tous les plongements élémentaires dans un groupe libre de type fini sont de cette forme, c'est-à-dire si un sous-groupe élémentaire d'un groupe libre est nécessairement un facteur libre. Un premier résultat dans cette direction peut être obtenu par des arguments simples :

Lemme 1.11 : *Soit H un sous-groupe élémentaire d'un groupe libre de type fini F . Alors H est un rétract de F .*

Démonstration. Notons que H est un groupe libre. On choisit $B_H = (h_1, h_2, \dots)$ une base (qui peut être infinie) pour H , et (a_1, a_2, \dots, a_n) une base pour F . Chacun des éléments h_i s'exprime par un mot w_i en les éléments a_j , on note $h_i = w_i(a_1, \dots, a_n)$.

On commence par montrer par contradiction que le rang de H est au plus n . Supposons que B_H a au moins $n + 1$ éléments : en particulier, H s'écrit comme un produit libre $H' * H''$, où H' est le sous-groupe librement engendré par h_1, \dots, h_{n+1} et H'' peut être trivial.

Considérons l'énoncé du premier ordre

$$\phi : \exists x_1 \dots x_n \bigwedge_{i=1}^{n+1} [h_i] = w_i(x_1, \dots, x_n)$$

C'est un énoncé dans \mathcal{L}_H , et il est satisfait par F : en effet, il suffit de prendre comme « solution » $x_j = a_j$. Comme H est un sous-groupe élémentaire de F , il satisfait également ϕ . Ceci implique l'existence d'éléments b_1, \dots, b_n de H tels que pour $1 \leq i \leq n + 1$, on a $h_i = w_i(b_1, \dots, b_n)$.

Soit B le sous-groupe de H engendré par b_1, \dots, b_n . Par le théorème de Kurosh, B hérite une décomposition en facteurs libres de la décomposition $H = H' * H''$, et l'un des facteurs de cette décomposition héritée est $B \cap H'$. Or pour $1 \leq i \leq n + 1$, on sait que $h_i = w_i(b_1, \dots, b_n)$ est dans B , donc $B \cap H' = H'$. Mais H' , qui est un groupe libre de rang $n + 1$, ne peut pas être un facteur libre de B qui est de rang au plus n : on a une contradiction. Le sous-groupe H est donc de rang au plus n .

On considère maintenant l'énoncé ϕ' donné par

$$\exists x_1 \dots x_n \bigwedge_{i=1}^k [h_i] = w_i(x_1, \dots, x_n),$$

où $k = \text{Card}(B_H)$. Cet énoncé est satisfait par F , donc par H , et comme précédemment on obtient des éléments b_1, \dots, b_n de H tels que pour tout $1 \leq i \leq k$, on a $h_i = w_i(b_1, \dots, b_n)$. Soit f le morphisme $G \rightarrow H$ défini par $f(a_j) = b_j$. On a $f(h_i) = f(w_i(a_1, \dots, a_n)) = w_i(b_1, \dots, b_n) = h_i$, donc f est une rétraction de F sur H . \square

Ceci n'est pas suffisant pour montrer que H doit être un facteur libre, mais on montrera le

Théorème A : *(Corollary 7.22) Un sous-groupe élémentaire plongé dans un groupe libre de type fini est un facteur libre.*

On obtiendra ce résultat comme corollaire du résultat principal de cette thèse, qui répond à la question un peu plus générale suivante :

Question 3: *Décrire les plongements élémentaires dans un groupe hyperbolique sans torsion.*

La description obtenue est donnée par le

Théorème B : *(Theorem 7.4) Soit G un groupe hyperbolique sans torsion. Soit H un groupe élémentaire plongé dans G . Alors G admet une structure de tour hyperbolique sur H .*

Les tours hyperboliques sont des structures définies par Sela (qui les appelle « hyperbolic ω -residually free towers »). Elles permettent de répondre à plusieurs questions de théorie des modèles sur les groupes libres et les groupes hyperboliques sans torsion. Ces structures font l'objet de la section suivante.

1.2 Tours hyperboliques

On donne la définition suivante :

Définition 1.12 : (tour hyperbolique) *Soient G un groupe et H un sous-groupe de G . On dira que G est une tour hyperbolique sur H s'il existe une suite finie $G = G^0 > G^1 > \dots > G^m > H$ de sous-groupes de G tels que :*

- *pour tout k dans $\llbracket 0, m-1 \rrbracket$, il existe une rétraction $r_k : G^k \rightarrow G^{k+1}$ telle que (G^k, G^{k+1}, r_k) est un étage hyperbolique.*
- *$G^m = H * F * S_1 * \dots * S_p$ où F est un groupe libre (éventuellement trivial), $p \geq 0$, et chaque S_i est le groupe fondamental d'une surface fermée de caractéristique d'Euler au plus -2 .*

On n'a pas défini la notion d'étage hyperbolique, ceci sera fait dans la définition 7.1. En attendant, en voici un exemple :

Exemple 1.13: Soit G un groupe, et soit $r : G \rightarrow G'$ une rétraction sur un sous-groupe de G . Supposons que G admet un scindement au-dessus d'un sous-groupe cyclique infini C de la forme $G = G' *_C S$, où S est le groupe fondamental d'une surface à une composante de bord, qui est soit un tore percé, soit de caractéristique d'Euler au plus -2 , et telle que le groupe fondamental de l'unique composante de bord est C . Si de plus l'image $r(S)$ de S par la rétraction est non abélienne, alors (G, G', r) est un étage de tour hyperbolique.

En général, S peut correspondre à une surface non connexe, à plusieurs composantes de bord. On supposera alors que l'image du groupe fondamental de chaque composante a une image non abélienne par la rétraction.

Exemple 1.14:

- Un groupe libre admet une structure de tour hyperbolique sur n'importe lequel de ses facteurs libres.
- Le groupe fondamental d'une surface fermée de caractéristique d'Euler au plus -2 admet une structure de tour hyperbolique sur 1. De même, un produit libre de groupes de surfaces fermées de caractéristique d'Euler au plus -2 est une tour hyperbolique sur 1, ou sur n'importe lequel de ses facteurs libres.
- Soit Σ une surface fermée de caractéristique d'Euler au plus -2 . Soit γ_0 une courbe fermée simple sur Σ qui sépare Σ en deux sous-surfaces Σ_0 et Σ_1 . On suppose que Σ_0 est soit un tore percé, soit de caractéristique d'Euler au plus -2 . Considérons le graphe de groupes à deux sommets de groupes $\pi_1(\Sigma)$ et $\pi_1(\Sigma_0)$ respectivement. Ces deux sommets sont joints par une arête de groupe cyclique infini qui s'envoie dans $\pi_1(\Sigma)$ isomorphiquement sur un groupe cyclique maximal correspondant à γ_0 , et dans $\pi_1(\Sigma_0)$ isomorphiquement sur un groupe de bord maximal. Alors, le groupe fondamental G de ce graphe de groupes est une tour hyperbolique sur $\pi_1(\Sigma)$. En effet, $\pi_1(\Sigma)$ contient un sous-groupe isomorphe à $\pi_1(\Sigma_0)$, l'application r qui est l'identité sur $\pi_1(\Sigma)$ et qui envoie $\pi_1(\Sigma_0)$ sur ce sous-groupe est bien définie, et fait de $(G, \pi_1(\Sigma), r)$ un étage de tour hyperbolique (voir figure 1.1).

La structure de tour hyperbolique apparaît dans plusieurs résultats de Sela. Par exemple, dans sa résolution du problème de Tarski, en plus de montrer que les groupes libres de type fini ont la même théorie élémentaire, Sela obtient une description des groupes de type fini qui ont la même théorie élémentaire qu'un groupe libre. Elle est donnée par le résultat suivant :

Théorème 1.15 : [Sel06, Proposition 6] *Soit G un groupe de type fini. Le groupe G est élémentairement équivalent à un groupe libre de type fini non abélien si et seulement s'il admet une structure de tour hyperbolique sur le groupe trivial.*




FIGURE 1.1 – Un exemple de tour hyperbolique.

Dans [Sel], qui fait suite à sa résolution du problème de Tarski, Sela généralise ses techniques des groupes libres aux groupes hyperboliques sans torsion. Il définit, pour tout groupe hyperbolique sans torsion Γ , un ensemble de sous-groupes deux à deux isomorphes, qu'il appelle des cœurs élémentaires (elementary cores), sur lesquels le groupe admet une structure de tour hyperbolique. Un cœur élémentaire H de Γ est tel que Γ n'admet pas de structure de tour hyperbolique sur un sous-groupe propre de H . Alors la classe d'isomorphisme des cœurs d'un groupe hyperbolique sans torsion Γ détermine sa classe d'équivalence élémentaire, comme l'exprime le

Théorème 1.16 : [Sel, Theorem 7.10] *Soit Γ un groupe hyperbolique non abélien et sans torsion. Si G est un groupe de type fini, G et Γ sont élémentairement équivalents si et seulement si G est hyperbolique sans torsion et les cœurs de G et Γ sont isomorphes.*

Sela montre également que si G est un groupe hyperbolique non abélien sans torsion qui n'est pas libre, le cœur de G est un sous-groupe élémentaire de G . Le Théorème B exprime que, comme le cœur, tout sous-groupe élémentaire forme la base d'une structure de tour hyperbolique pour G .

1.3 Structure de la preuve du théorème B

Soit H un sous-groupe élémentaire d'un groupe hyperbolique sans torsion G . Pour montrer que G admet une structure de tour hyperbolique sur H , il faut commencer par trouver l'étage supérieur de la tour : on veut trouver une rétraction r de G dans un sous-groupe G' telle que (G, G', r) est un étage hyperbolique.

Pour ce faire, on utilise un résultat, implicite dans la preuve de la proposition 6 de [Sel06], qui nous permet de construire une telle rétraction à partir d'un morphisme $G \rightarrow G$ qui respecte certaines propriétés d'une décomposition en graphe de groupes Λ de G . Cette décomposition doit satisfaire certaines propriétés d'acylindricité, et certains de ses groupes de sommets sont des groupes fondamentaux de surfaces à bord dont les sous-groupes de bord sont exactement les groupes d'arêtes adjacents. De tels sommets sont appelés sommets de type surface. Une décomposition qui satisfait ces hypothèses sera appelée décomposition de type JSJ, puisqu'en pratique, on considérera la plupart du temps soit la décomposition JSJ (voir [RS97]), soit la décomposition JSJ relative à un sous-groupe.

Le résultat qu'on utilise est donné par la proposition suivante, qui apparaîtra sous une forme légèrement différente dans la proposition 7.15.

Proposition C : *Soit G un groupe hyperbolique non abélien sans torsion. Soit Λ une décomposition de type JSJ de G . On suppose qu'il existe un morphisme $f : G \rightarrow G$ non injectif tel*

que

- si R est un groupe de sommet de Λ qui n'est pas de type surface, la restriction de f à R est une conjugaison par un élément g_R de G ;
- si S est un groupe de sommet de type surface de Λ , $f(S)$ n'est pas abélien.

Alors il existe une rétraction r de G sur un sous-groupe G' , telle que (G, G', r) est un étage hyperbolique. De plus, si R_0 est un groupe de sommet qui n'est pas de type surface, on peut choisir r de telle manière que $R_0 \leq G'$.

Pour trouver l'étage supérieur d'une structure de tour hyperbolique pour G , il nous suffit donc de montrer qu'un morphisme $f : G \rightarrow G$ satisfaisant les hypothèses de la proposition C existe. C'est là qu'on utilisera la logique du premier ordre. Nous allons maintenant essayer d'indiquer quelques éléments des deux étapes principales de la preuve : la preuve de la proposition C, et la construction d'un morphisme qui satisfait les conditions de la proposition C.

1.3.1 Construction du morphisme f

Soit donc G un groupe hyperbolique non cyclique et sans torsion, et soit H un sous-groupe propre de G dont l'inclusion dans G est élémentaire.

Supposons pour simplifier que G est librement indécomposable par rapport à H . On considère Λ la décomposition JSJ de G par rapport à H . On supposera aussi que H est finiment engendré : attention, ceci n'est pas nécessairement le cas à priori, et l'argument qu'on utilise dans le cas des groupes libres ne se généralise pas ici. En fait, on obtiendra que H est finiment engendré seulement comme conséquence du Théorème B.

Si Λ est triviale, on n'a aucun espoir de trouver un morphisme f qui satisfait les hypothèses de la proposition C : en effet, G lui-même est un groupe de sommet qui n'est pas de type surface donc f est simplement une conjugaison. Mais alors f est nécessairement injective. Heureusement, on a le

Lemme 1.17 : *La décomposition Λ n'est pas triviale.*

On utilisera l'existence d'un ensemble de factorisation pour $\text{Hom}_H(G, G)$, décrit par le résultat suivant :

Théorème D : *(cas particulier du Theorem 6.29) Soit G un groupe hyperbolique non abélien, sans torsion et librement indécomposable par rapport à un sous-groupe H . Il existe un ensemble fini de quotients propres $\{\eta_1 : G \rightarrow L_1, \dots, \eta_k : G \rightarrow L_k\}$, appelé ensemble de factorisation pour $\text{Hom}_H(G, G)$, tel que tout morphisme non injectif $G \rightarrow G$ qui fixe H se factorise par l'un de ces quotients après précomposition par un automorphisme modulaire de G relativement à H .*

Ce résultat est bien évidemment à rapprocher de l'existence d'un ensemble de factorisation pour $\text{Hom}(A, \mathbb{F}_n)$, l'ensemble des morphismes d'un groupe A de type fini dans un groupe libre (voir proposition 6.9). Cela donne une preuve de ce résultat dans [Sel01], qu'il généralise ensuite dans [Sel] à la preuve de l'existence d'un ensemble de factorisation pour $\text{Hom}(A, \Gamma)$, où Γ est un groupe hyperbolique sans torsion (voir proposition 6.19). On montrera une version relative de ce résultat (proposition 6.29), c'est-à-dire l'existence d'un ensemble de factorisation pour $\text{Hom}_H(A, \Gamma)$ où H se plonge dans A et dans Γ . Le théorème D est un cas particulier de ce résultat dans le cas où $A = \Gamma = G$.

Démonstration du lemme 1.17. Fixons quelques notations. Soit h_1, \dots, h_n une partie génératrice pour H . On choisit également une présentation finie de G donnée par $G = \langle \mathbf{g} \mid \Sigma_G(\mathbf{g}) \rangle$ où Σ_G dénote un ensemble fini de mots en les éléments de \mathbf{g} . Chaque h_j est représenté par un mot $h_j(\mathbf{g})$.

Si Λ est triviale, le groupe modulaire de G relativement à H est trivial également, donc le théorème D nous donne un ensemble fini de quotients propres η_i de G , tel que tout morphisme non injectif de G dans G qui fixe H se factorise par l'un des η_i . On choisit pour chaque i un élément non trivial v_i de $\text{Ker}(\eta_i)$. Chaque v_i est représenté par un mot $v_i(\mathbf{g})$.

Notons maintenant que si ϕ est un morphisme de G dans H qui fixe H , il ne peut pas être injectif puisque H est un sous-groupe propre de G . Par ailleurs, comme c'est également un morphisme de G dans G , il doit se factoriser par l'un des quotients η_i .

L'ensemble des morphismes $G \rightarrow H$ est en bijection avec l'ensemble des solutions de l'équation $\Sigma_G(\mathbf{g}) = 1$ dans H : à une solution \mathbf{x} est associé le morphisme $\phi_{\mathbf{x}}$ qui envoie \mathbf{g} sur \mathbf{x} . L'image par $\phi_{\mathbf{x}}$ d'un élément représenté par un mot $w(\mathbf{g})$ est alors représentée par le mot $w(\mathbf{x})$. On voit donc que le morphisme $\phi_{\mathbf{x}}$ fixe H si et seulement si on a $h_i = h_i(\mathbf{x})$ pour tout i .

Comme tout morphisme de G dans H qui fixe H se factorise par l'un des quotients η_i , l'énoncé du premier ordre sur \mathcal{L}_H

$$\forall \mathbf{x} [\Sigma_G(\mathbf{x}) = 1 \wedge \bigwedge_{i=1}^n [h_i] = h_i(\mathbf{x})] \rightarrow \bigvee_i^r v_i(\mathbf{x}) = 1$$

est satisfait par H .

Puisque H est plongé élémentairement dans G , cet énoncé doit également être satisfait par G . Mais prenons dans G la solution tautologique $\mathbf{x} = \mathbf{g}$. Elle satisfait l'équation $\Sigma_G(\mathbf{g}) = 1$, et on a par définition $h_i = h_i(\mathbf{g})$ pour tout i . Cependant, aucun des mots $v_i(\mathbf{g})$ ne représente l'élément neutre. C'est une contradiction : la décomposition JSJ de G relativement à H n'est pas triviale. \square

De la même manière, si Λ ne contient pas de sommets de type surface, un morphisme f qui satisfait les hypothèses de la proposition C est un isomorphisme. Là encore, on peut montrer comme dans la preuve du lemme précédent que ce cas ne se produit pas.

Pour le cas général, on utilisera de même l'existence d'un ensemble de factorisation pour trouver un énoncé du premier ordre dans \mathcal{L}_H satisfait par H . Cependant, lorsque le groupe modulaire est suffisamment complexe, c'est-à-dire quand la décomposition JSJ comporte des sommets de type surface, il est impossible d'exprimer l'existence d'un ensemble de factorisation par une formule du premier ordre. L'énoncé qu'on considère exprime alors une affirmation plus faible, et son interprétation dans G nous permettra de trouver un morphisme $f : G \rightarrow G$ qui satisfait les conditions de la proposition C.

1.3.2 Preuve de la proposition C

On considère deux possibilités simples pour la décomposition Λ .

Exemple 1.18: Supposons que Λ est un graphe à deux sommets de groupes A et B qui ne sont pas de type surface et à une arête joignant ces deux sommets. Un morphisme $f : G \rightarrow G$ dont la restriction à A et à B est une conjugaison par des éléments g_A et g_B respectivement est un automorphisme. Il est donc forcément injectif, et la proposition est trivialement vraie.

Exemple 1.19: Supposons maintenant que Λ est un graphe sur deux sommets v_A et v_S et une arête les joignant. On suppose de plus que seul v_S est de type surface, et on dénote A et S les groupes de sommets de A et B respectivement. Soit $f : G \rightarrow G$ un morphisme qui satisfait les conditions de la proposition C. On suppose de plus qu'aucun élément correspondant à une courbe fermée simple sur la surface correspondant à S n'est dans le noyau de f . Quitte à conjuguer f , on peut supposer que f est l'identité sur A .

On considère l'image de S par f . Si $f(S) \leq A$, alors f est une rétraction de G dans A , et on voit ensuite aisément que (G, A, f) est un étage de tour hyperbolique.

Si $f(S) \leq S$, en utilisant le fait que f ne tue aucun élément correspondant à une courbe fermée simple, on peut montrer que $f(S)$ doit être un sous-groupe d'indice fini de S . Mais le rang de $f(S)$ est au plus égal au rang de S , or le rang d'un sous-groupe d'indice fini dans un groupe libre de rang k est de rang au moins k , avec égalité si et seulement si l'indice est 1. On en déduit que $f(S) = S$. Comme les groupes libres sont hopfien, f restreint à S est un isomorphisme. Donc f est un isomorphisme, ce qui contredit sa non-injectivité.

Pour traiter le cas général, on note que S agit via f sur l'arbre simplicial correspondant à Λ : il hérite donc d'une décomposition en graphe de groupes, dont on peut montrer qu'elle est duale à un ensemble de courbes fermées simples sur la surface Σ correspondante. Ces courbes divisent Σ en un nombre fini de sous-surfaces dont les groupes fondamentaux ont une image par f elliptique pour Λ . Si une telle sous-surface Σ_0 a pour groupe fondamental S_0 , et si $f(S_0)$ est un sous-groupe non abélien de S , on peut appliquer un argument similaire à celui du paragraphe précédent pour voir que Σ_0 doit être au moins aussi complexe que Σ (pour une notion de complexité un peu plus précise que le rang du groupe fondamental). Ceci n'est pas possible si Σ_0 est une sous-surface propre de Σ . On peut donc montrer que les groupes fondamentaux de toutes les sous-surfaces obtenues sont envoyés par f dans un conjugué de A . Ceci permet finalement de voir que f est la rétraction cherchée.

1.4 Contenu de la thèse

Dans le chapitre 3 sont exposés quelques rappels sur la théorie de Bass-Serre, qui décrit les actions de groupe sur des arbres simpliciaux. On rappellera aussi quelques propriétés élémentaires des espaces métriques hyperboliques, et on définira la topologie de Gromov-Hausdorff. Les chapitres 4, 5 et 6 ont pour thème l'argument du raccourcissement et certaines de ses conséquences. Dans le chapitre 4, on donne diverses variantes de l'argument du raccourcissement, et on présente une preuve de certains résultats de raccourcissement sur des suites de morphismes. Pour ce faire, on utilise les résultats de raccourcissement sur des suites d'actions, dont la preuve est l'objet du chapitre 5. Ces résultats nous permettent de montrer tout d'abord certaines propriétés de type co-Hopf pour les groupes hyperboliques, puis l'existence d'ensembles de factorisations dans le chapitre 6 : là encore, on énonce les diverses versions de ce résultat (pour les morphismes vers les groupes libres, vers les groupes hyperboliques sans torsion, relativement à un sous-groupe), qui proviennent de diverses formes de l'argument du raccourcissement.

Dans le reste de la thèse, on donne la preuve du Théorème B : le chapitre 7 expose le résultat qu'on prouve en utilisant la proposition C. Enfin, les chapitres 8, 9 et 10 sont consacrés à la preuve de la proposition C.

Chapter 2

Introduction in English

In the last decades, the apparition of geometric group theory has led to the resolution of many problems in groups theory. The introduction of typically geometric concepts such as curvature, or geodesics, has allowed to tackle questions that the traditional combinatorial or purely algebraic approach had left unsolved. Representing a group by a metric space allows to define new notions such as quasi-isometry between groups, and new classes of groups like that of hyperbolic or CAT(0) groups. The geometric properties of the space associated to a group are of independent interest, but have also proved to bear strong relation to the classical algebraic properties of the group.

More recently, such a geometrical approach has proved particularly fruitful when applied to questions borrowed from model theory. This new interaction finds its source mainly in the work of Sela on Tarski's problem, in [Sel01]-[Sel] (see also Kharlampovich and Myasnikov's approach in [KM06]). The results of Sela make an extensive use of Bass-Serre theory about actions on simplicial trees, of the JSJ theory, which describes all such possible actions, and of Rips theory, which analyses actions on real trees. Here again, the results highlight a deep connection between the geometric properties of a group, and its first-order theory. For example, Sela shows that a group which has the same first-order theory as a torsion-free hyperbolic group must be torsion-free hyperbolic itself.

This geometric approach has thus been useful to solve several difficult model-theoretical questions, but has also considerably enriched the tools of geometric group theory. Moreover, some problems which at first might seem unrelated to both areas have been resolved along the way, such as the study of ω -residually free groups, or the resolution of equations over free groups. For example, Sela obtains in [Sel01] a new proof of the existence of a Makanin-Razborov diagram (the original result was proved by Razborov in [Raz85], and generalised by Kharlampovich and Myasnikov in [KM98]). Such a diagram classifies homomorphisms from a given finitely generated group into a free group. Because the proof is essentially geometric, it can be generalised to torsion-free hyperbolic groups, whose geometry is close to that of free groups, so Sela gets in [Sel] a Makanin-Razborov diagram for morphisms into a torsion-free hyperbolic group. Groves generalises this result further to groups that are hyperbolic with respect to a collection of free abelian groups in [Gro05].

One of the basic notions of model theory is that of an elementary embedding, which describes how a structure embeds in another in a way which makes them indistinguishable for first-order theory. The main purpose of this thesis is to study elementary embeddings in torsion-free hyperbolic group.

2.1 Elementary embeddings

The **language of groups** is the following set of symbols

$$\mathcal{L} = \{=, (,), \neg, \vee, \wedge, \forall, \exists, 1, *, ^{-1}\} \cup V$$

where V is an infinite countable set of variables. Recall that the symbol \vee represents disjunction ("or"), the symbol \wedge conjunction ("and"), and the symbol \neg negation ("not"). The symbol 1 represents the unit element of the group, $*$ is the multiplication (but we will mostly represent product simply by concatenation), and $^{-1}$ denotes the inverse. A **first-order formula** (or **elementary formula**) in the language \mathcal{L} is a finite sequence of elements of \mathcal{L} which constitutes a "grammatically correct" mathematical formula. In the sequel, we will often use usual mathematical symbols to represent a finite set of elements of \mathcal{L} , such as for example the symbol \rightarrow , where $A \rightarrow B$ represents $B \vee \neg A$.

A variable x which appears in a first-order formula is **free** if neither $\forall x$ nor $\exists x$ appears before it in the formula. A first-order formula ϕ is said to be **closed** (we will also say that ϕ is a **sentence**) if none of the variables which appear in ϕ are free. A group G **satisfies** a sentence ϕ of the language \mathcal{L} if the interpretation of the formula holds in G . We denote this by $G \models \phi$.

Example 2.1: If ϕ is the formula $\forall x \forall y x * y * x^{-1} * y^{-1} = 1$, a group G satisfies ϕ if and only if it is abelian.

Definition 2.2: (elementary theory) *Let G be a group. The elementary theory of G in \mathcal{L} is the set of closed first-order formulas over \mathcal{L} satisfied by G .*

It is important to note that quantification is allowed only on elements of the group. In particular, it is not allowed on subsets of the group, nor on integers.

Example 2.3: To express the fact that a group is torsion-free, we might want to write the following formula

$$\forall x (x \neq 1) \rightarrow \bigwedge_{n=1}^{\infty} (x^n \neq 1)$$

However, this is **not** a first-order formula, since if we rewrote it using only symbols of \mathcal{L} (without short cuts), we would get an infinite formula. Similarly, the formula

$$\forall x (x \neq 1) \rightarrow \forall n \in \mathbb{N} (x^n \neq 1)$$

is **not** a first-order formula, since we quantify on integers.

Definition 2.4: (elementary equivalent) *Two groups G and G' are elementary equivalent if they have the same elementary theory in the language of groups. We denote this by $G \equiv G'$.*

Example 2.5: Let G and G' two groups for which $G \equiv G'$.

- If G is abelian, so is G' .
- If G is finite, so is G' , and they have the same cardinality. In fact, they are isomorphic: the multiplication table of G can be expressed by a first-order formula, which is satisfied by G' .
- If G is torsion-free, so is G' . Indeed, even though 'being torsion-free' cannot be expressed by **one** first-order formula, it can be expressed by the following infinite family of sentences:

$$\{\forall x [(x \neq 1) \rightarrow (x^n \neq 1)]\}_{n \in \mathbb{N} - \{0\}}.$$

If G is torsion-free, it satisfies each one of these formulas, hence so does G' .

Example 2.6: The groups \mathbb{Z} and \mathbb{Z}^2 are not elementary equivalent. Indeed, \mathbb{Z} satisfies the following sentence

$$\exists x \forall y \exists z (y = z^2) \vee (y = z^2 x)$$

which expresses that in \mathbb{Z} , an element is either even or odd. Clearly, \mathbb{Z}^2 does not satisfy this sentence. It can be showed in this way that $\mathbb{Z}^k \equiv \mathbb{Z}^l$ if and only if $k = l$.

We can now state the following problem:

Question 4: *Suppose that $1 < m < n$. Are free groups of rank m and n elementary equivalent?*

This question was asked by the logician Alfred Tarski around 1945, and is known as Tarski's problem. Sela answered it positively in [Sel06]. Kharlampovich and Myasnikov have another approach to this problem (see [KM06]). Sela also gives a characterisation of all finitely generated groups which are elementary equivalent to non-abelian free groups (see Theorem 2.15). The connection with geometry is striking in the following result, which is a corollary of this characterisation:

Theorem 2.7: *The fundamental group of a closed surface whose Euler characteristic is at most -2 is elementary equivalent to a non-abelian free group.*

It is natural to study in the context of free groups other classical notions of model theory, such as that of elementary embedding.

Definition 2.8: (elementary embedding) *Let G be a group, and let H be a subgroup of G . We denote by \mathcal{L}_H the language of groups \mathcal{L} to which have been added for each element h of H a new constant $[h]$. We say that the embedding $H \subseteq G$ is elementary, or that H is an elementary subgroup of G if for any first-order sentence ϕ in the language \mathcal{L}_H , the subgroup H satisfies ϕ if and only if G satisfies ϕ . We denote this by $H \preceq G$.*

Note that this definition is equivalent to saying that H and G are elementary equivalent in the language \mathcal{L}_H , and thus implies standard elementary equivalence (in the language \mathcal{L}).

Example 2.9: Let G be a group, let H be a subgroup of G . Let h be an element of H . Consider the following sentence

$$\phi_h : \forall x [[h], x] = 1$$

It is a first-order sentence in the language \mathcal{L}_H . The group H (respectively G) satisfies ϕ_h if and only if h lies in the centre $Z(H)$ of H (respectively $Z(G)$ of G).

In particular, if $H \preceq G$ we see that $h \in Z(H)$ if and only if $h \in Z(G)$, we thus have $Z(H) = H \cap Z(G)$.

When studying the first-order theory of free groups, the following question comes up naturally:

Question 5: *Describe elementary embeddings in free groups.*

In his proof of the elementary equivalence of free groups, Sela shows in fact the

Theorem 2.10: (Theorem 4 of [Sel06]) *Let $i : \mathbb{F}_k \rightarrow \mathbb{F}_n$ be the canonical embedding of a free group of rank k in the free group of rank n for $2 \leq k \leq n$. Then i is an elementary embedding.*

It is thus natural to ask if all the elementary embeddings in a free groups are of this type, that is, whether an elementary subgroup of a free group is necessarily a free factor. A first step in this direction is:

Lemma 2.11: *Let H be an elementary subgroup of a finitely generated free group F . Then H is a retract of F .*

Proof. Note that H is a free group. We chose $B_H = (h_1, h_2, \dots)$ a basis for H (it might be infinite), and (a_1, a_2, \dots, a_n) a basis for F . Each element h_i can be expressed by a word w_i in the elements a_j , we write $h_i = w_i(a_1, \dots, a_n)$.

Let us first see by contradiction that the rank of H is at most n . Suppose that B_H has at least $n + 1$ elements: in particular, H can be written as a free product $H' * H''$, where H' is the subgroup freely generated by h_1, \dots, h_{n+1} and H'' is possibly trivial. Consider the following first-order sentence

$$\phi : \exists x_1 \dots x_n \bigwedge_{i=1}^{n+1} [h_i] = w_i(x_1, \dots, x_n).$$

It is a sentence of \mathcal{L}_H which is satisfied by F : indeed, we can take $x_j = a_j$ as a "solution". Since H is an elementary subgroup of F , it also satisfies ϕ . This implies that there exist elements b_1, \dots, b_n of H such that for any $1 \leq i \leq n+1$, we have $h_i = w_i(b_1, \dots, b_n)$. Let B be the subgroup of H generated by b_1, \dots, b_n . By Kurosh's theorem, B inherits a free product decomposition from $H = H' * H''$, in which one of the factors is $B \cap H'$. But for $1 \leq i \leq n+1$, we have $h_i = w_i(b_1, \dots, b_n)$ so $h_i \in B$. Thus $B \cap H' = H'$ is a free group of rank $n+1$, so it cannot be a free factor of B whose rank is at most n : we get a contradiction. The subgroup H has rank at most n .

We now consider the sentence ϕ' given by

$$\exists x_1 \dots x_n \bigwedge_{i=1}^k [h_i] = w_i(x_1, \dots, x_n).$$

where $k = \text{Card}(B_H)$. It is satisfied by F , thus it is satisfied by H and we get elements b_1, \dots, b_n of H such that for $1 \leq i \leq k$, we have $h_i = w_i(b_1, \dots, b_n)$. Let f be the morphism $G \rightarrow H$ defined by $f(a_j) = b_j$. We get $f(h_i) = f(w_i(a_1, \dots, a_n)) = w_i(b_1, \dots, b_n) = h_i$, thus f is a retraction from F to H . \square

This is not enough to show that H must be a free factor, but we will see:

Theorem A: (Corollary 7.22) *An elementary subgroup of a finitely generated free group is a free factor.*

This will be a corollary of the main result of the thesis, which answers the following question:

Question 6: *Describe elementary embeddings in torsion-free hyperbolic groups.*

The description we obtain is given by

Theorem B: (Theorem 7.4) *Let G be a torsion-free hyperbolic group. Let H be a subgroup elementarily embedded in G . Then G has a structure of hyperbolic tower over H .*

Hyperbolic towers are structures defined by Sela (who calls them 'hyperbolic ω -residually free towers'). They appear in the answer to several questions about the first-order theory of free and hyperbolic groups. They are the subject of the following section.

2.2 Hyperbolic towers

We give the following definition:

Definition 2.12: (hyperbolic tower) *Let G be a group, and let H be a subgroup of G . We say that G is a hyperbolic tower over H if there is a finite sequence $G = G^0 > G^1 > \dots > G^m > H$ of subgroups of G such that:*

- *for any k in $[0, m-1]$, there is a retraction $r_k : G^k \rightarrow G^{k+1}$ such that (G^k, G^{k+1}, r_k) is a hyperbolic floor.*
- *$G^m = H * F * S_1 * \dots * S_p$ where F is a (possibly trivial) free group, $p \geq 0$, and each S_i is the fundamental group of a closed surface of Euler characteristic at most -2 .*




Figure 2.1: An example of hyperbolic tower.

We have not defined hyperbolic floors yet, this will be done in definition 7.1. In the meantime, let us give an example:

Example 2.13: Let G be a group, and let $r : G \rightarrow G'$ be a retraction onto a subgroup of G . Suppose that G has a splitting over an infinite cyclic subgroup C of the form $G = G' *_C S$, where S is the fundamental group of a surface with exactly one boundary component, which either is a punctured torus, or has Euler characteristic at most -2 , and such that the fundamental group of the unique boundary component is C . If moreover the image $r(S)$ of S by the retraction is non-abelian, then (G, G', r) is a hyperbolic floor.

In general, S might correspond to a disconnected surface, with several boundary components. We will then assume that the image of the fundamental group of each connected component has non-abelian image by the retraction.

Example 2.14:

- A free group has a structure of hyperbolic tower over each of its free factors.
- The fundamental group of a closed surface of Euler characteristic at most -2 has a structure of hyperbolic tower over 1. Similarly, a free product of fundamental groups of surfaces of Euler characteristic at most -2 is a hyperbolic tower over 1, and over each of its free factors.
- Let Σ be a closed surface of Euler characteristic at most -2 . Let γ_0 be a simple closed curve on Σ which separates Σ into two subsurfaces Σ_0 and Σ_1 . We assume that Σ_0 is either a punctured torus, or of Euler characteristic at most -2 . Consider the graph of group on two vertices of groups $\pi_1(\Sigma)$ and $\pi_1(\Sigma_0)$ respectively, joined by an edge of infinite cyclic edge group, which injects in $\pi_1(\Sigma)$ isomorphically onto a maximal cyclic subgroup corresponding to γ_0 , and in $\pi_1(\Sigma_0)$ isomorphically on a maximal boundary subgroup. Then, the fundamental group G of this graph of groups is a hyperbolic tower over $\pi_1(\Sigma)$. Indeed, $\pi_1(\Sigma)$ contains a subgroup isomorphic to $\pi_1(\Sigma_0)$: the map r which restricts to the identity on $\pi_1(\Sigma)$ and which sends $\pi_1(\Sigma_0)$ on this subgroup is well-defined, and makes $(G, \pi_1(\Sigma), r)$ a hyperbolic floor (see figure 2.1).

Hyperbolic towers appear in several results of Sela. For example, in his solution to Tarski's problem, as well as showing that finitely generated free groups are all elementary equivalent, Sela gives a description of finitely generated groups which have the same elementary theory as a free group. It is given by the following result:

Theorem 2.15: (*Proposition 6 [Sel06]*) *Let G be a finitely generated group. The group G is elementary equivalent to a non-abelian finitely generated free group if and only if it admits a*

structure of hyperbolic tower over the trivial group.

In a paper which follows his resolution of Tarski's [Sel], Sela generalises his techniques from free groups to torsion-free hyperbolic groups. Given a torsion-free hyperbolic group Γ , he defines subgroups of Γ called elementary cores, which are all isomorphic, and over which Γ admits a structure of hyperbolic tower. A core H of Γ is such that Γ is not a hyperbolic tower over any proper subgroup of H . Then, the isomorphism class of the cores of a torsion-free hyperbolic group Γ determines its elementary equivalence class, so that we have

Theorem 2.16: *(Theorem 7.10 in [Sel]) Let Γ be a non-abelian, torsion-free hyperbolic group. If G is a finitely generated group, G and Γ are elementary equivalent if and only if G is torsion-free hyperbolic and the cores of G and Γ are isomorphic.*

Sela shows also that if Γ is a non-abelian torsion-free hyperbolic group which is not free, the core of Γ is an elementary subgroup of Γ . Theorem B says that, as the core, any elementary subgroup is the basis of a hyperbolic tower structure for Γ .

2.3 Structure of the proof of theorem B

Let H be an elementary subgroup of a torsion-free hyperbolic group G . To show that G has a structure of hyperbolic tower over H , we must first find the top floor of the tower. In other words, we want to show that there is a retraction r from G to a subgroup G' such that (G, G', r) is a hyperbolic floor.

To do this, we use a result which is implicit in the proof of proposition 6 of [Sel06], and which enables us to build such a retraction from a morphism $G \rightarrow G$ which respects some properties of a graph of groups decomposition Λ of G . This decomposition must satisfy some conditions of acylindricity, and some of its vertices are fundamental groups of surfaces with boundary whose boundary subgroups are exactly the adjacent edge groups. Such vertices are called surface type vertices. A decomposition which satisfies these hypotheses will be called a JSJ-like decomposition, since in fact most of the decompositions of this type that we consider will be either JSJ decompositions (see [RS97]), or JSJ decompositions relative to a subgroup.

The result we use is given by the following proposition, which will appear in a slightly different form in proposition 7.15.

Proposition C: *Let G be a non-abelian torsion-free hyperbolic group. Let Λ be a JSJ-like decomposition of G . Suppose that there exists a non-injective morphism $f : G \rightarrow G$ such that*

- *if R is a vertex group of Λ which is not of surface type, the restriction of f to R is a conjugation by an element g_R of G ;*
- *if S is a vertex group of Λ which is of surface type, $f(S)$ is not abelian*

Then there exists a retraction r from G onto a subgroup G' , such that (G, G', r) is a hyperbolic floor. Moreover, if R_0 is a vertex group which is not of surface type, we can choose r such that $R_0 \leq G'$.

Thus, to find the top floor of a hyperbolic tower structure for G , it is enough to show that there exists a morphism which satisfies the hypotheses of proposition C. This is where we use first-order logic. We will now give a few elements of the two main steps of the proof of Theorem B: the proof of proposition C, and the construction of a morphism satisfying the conditions of proposition C.

2.3.1 Construction of the morphism f

Let thus G be a non-abelian torsion-free hyperbolic group, and let H be a proper subgroup of G whose embedding in G is elementary.

Let us assume for the sake of simplicity that G is freely indecomposable with respect to H , and consider Λ the JSJ decomposition of G with respect to H . We also assume that H is finitely generated: this is not necessarily the case a priori, and the argument we gave in the case of free groups does not generalise here. In fact, we will only get that H is finitely generated as a consequence of Theorem B.

If Λ is trivial, there is no hope to find a morphism f which satisfies the hypotheses of proposition C: indeed, G itself is a vertex group which is not of surface type, so f is just a conjugation. But this means f has to be injective. Luckily, we have

Lemma 2.17: *The decomposition Λ is not trivial.*

To prove this, we use the existence of a factor set for $\text{Hom}_H(G, G)$, described in the following result:

Theorem D: *(particular case of Theorem 6.29) Let G be a non-abelian torsion-free hyperbolic group which is freely indecomposable with respect to a subgroup H . There is a finite set of proper quotients $\eta_1 : G \rightarrow L_1, \dots, \eta_k : G \rightarrow L_k$, called a factor set for $\text{Hom}_H(G, G)$, such that any non-injective morphism $G \rightarrow G$ fixing H factorises through one of these quotients after precomposition by a modular automorphism of G relative to H .*

This result is of course related to the existence of a factor set for $\text{Hom}(A, \mathbb{F}_n)$, the set of morphisms from a finitely generated group A to a free group (see proposition 6.9). Sela gives a proof of this result in [Sel01], that he then generalises in [Sel] to the proof of the existence of a factor set for $\text{Hom}(A, \Gamma)$, where Γ is a torsion-free hyperbolic group (see proposition 6.19). We will show a relative version of this result (proposition 6.29), that is we will show the existence of a factor set for $\text{Hom}_H(A, \Gamma)$ where H embeds both in A and in Γ . Theorem D is the particular case of this result when $A = \Gamma = G$.

Proof. Let us fix some notations. Let h_1, \dots, h_n be a generating set for H . We also choose a finite presentation of G given by $G = \langle \mathbf{g} \mid \Sigma_G(\mathbf{g}) \rangle$ where Σ_G denotes a finite set of words in the elements of \mathbf{g} . Each h_j is represented by a word $h_j(\mathbf{g})$ in the elements \mathbf{g} .

If Λ is trivial, the modular group of G relatively to H is trivial, so Theorem D gives us a finite set of proper quotients η_i of G , such that any non-injective morphism from G to G which fixes H factorises through one of the maps η_i . We pick for each i a non-trivial element v_i of $\text{Ker}(\eta_i)$. Each v_i is represented by a word $v_i(\mathbf{g})$.

Let us restrict ourselves to the set of morphisms $G \rightarrow H$: it is in bijection with the set of solutions of the equation $\Sigma_G(\mathbf{g}) = 1$ in H . To a solution \mathbf{x} is associated the morphism $\phi_{\mathbf{x}}$ which sends \mathbf{g} to \mathbf{x} . The image by $\phi_{\mathbf{x}}$ of an element represented by the word $w(\mathbf{g})$ is then represented by $w(\mathbf{x})$. Thus we see that the morphism $\phi_{\mathbf{x}}$ fixes H if and only if we have $h_i = h_i(\mathbf{x})$ for all i .

Now remark that if ϕ is a morphism from G to H which fixes H , it cannot be injective since H is a proper subgroup of G . On the other hand, it is also a morphism from G to G , so it must factorise through one of the quotients η_i and for some index i we have $v_i(\mathbf{x}) = 1$.

This tells us that the following sentence

$$\forall \mathbf{x} \left[\Sigma_G(\mathbf{x}) = 1 \wedge \bigwedge_{i=1}^n [h_i] = h_i(\mathbf{x}) \right] \rightarrow \bigvee_i^r v_i(\mathbf{x}) = 1$$

is satisfied by H .

Since H is elementarily embedded in G , this sentence must also be satisfied by G . But take in G the tautological solution $\mathbf{x} = \mathbf{g}$. It satisfies the equation $\Sigma_G(\mathbf{g}) = 1$, and by definition, we

have $h_i = h_i(\mathbf{g})$ for all i . However, none of the words $v_i(\mathbf{g})$ represents the trivial word. This is a contradiction, so the JSJ decomposition of G relative to H isn't trivial. \square

Similarly, if Λ does not contain surface type vertices, a morphism f which satisfies the hypotheses of proposition C is an isomorphism. We can show as in the previous lemma that this case does not occur.

For the general case, we will also use the existence of a factor set to find a first-order sentence on \mathcal{L}_H satisfied by H . However, when the modular group is complex enough, namely when the JSJ decomposition contains surface type vertices, it is impossible to express the existence of a factor set by a first-order formula. The sentence we then consider expresses something weaker, and its interpretation in G will give us a morphism $f : G \rightarrow G$ which satisfies the hypotheses of proposition C.

2.3.2 Proof of Proposition C

We consider two simple cases for the decomposition Λ .

Example 2.18: Suppose that Λ is a graph of groups on two vertices which are not of surface type, whose groups we denote by A and B , and with a single edge joining the two vertices. A morphism $f : G \rightarrow G$ whose restriction to A and to B is a conjugation by elements g_A and g_B respectively is an automorphism. The proposition is trivially true.

Example 2.19: Suppose now that Λ is a graph on two vertices v_A and v_S of groups A and S , with v_S of surface type, and a single edge joining the two vertices. Let $f : G \rightarrow G$ be a morphism which satisfies the conditions of proposition C. We assume moreover that no element corresponding to a simple closed curve on the surface corresponding to S is in the kernel of f . Up to conjugating f , we may assume that f is the identity on A .

Consider the image of S by f . If $f(S) \leq A$, f is a retraction from G to A , and we see easily that (G, A, f) is a hyperbolic floor.

If $f(S) \leq S$, using the fact that f does not kill elements corresponding to simple closed curves, we can show that $f(S)$ must be a finite index subgroup of S . But the rank of $f(S)$ is at most equal to the rank of S , and the rank of a finite index subgroup in a free group of rank k has rank at least k , with equality if and only if the index is 1. We deduce that $f(S) = S$. Since free groups are Hopfian, f restricted to S is an isomorphism. Thus f is an isomorphism, which contradicts its non-injectivity.

To deal with the general case, note that S acts via f on the simplicial tree corresponding to Λ : it inherits a graph of groups decomposition. We can show that this decomposition is dual to a set of simple closed curves on the surface Σ . These curves divide Σ into a finite number of subsurfaces whose fundamental groups have elliptic image for Λ . If such a subsurface Σ_0 has fundamental group S_0 , and if $f(S_0)$ is a non-abelian subgroup of S , we can use an argument similar to the one above to see that Σ_0 must be at least as complex as Σ (for a notion of complexity which is slightly more complicated than the rank of the fundamental group). This is not possible if Σ_0 is a proper subsurface of Σ . We can thus show that the fundamental groups of all the subsurfaces are sent to a conjugate of A by f . This finally enables us to see that f is the retraction we were looking for.

2.4 Content of the thesis

The first half of the thesis revolves around the shortening argument and some of its consequences. We start in chapter 3 by recalling some basic notions about graphs of groups, hyperbolic metric spaces and Gromov-Hausdorff topology. In chapter 4, we state various versions of the shortening

argument, and we present a proof of the shortening for a sequence of morphisms in the classical and the relative case. To do so, we use the shortening result for a sequence of actions, whose proof is the object of chapter 5. This enables us to show first some properties of Co-Hopf type for hyperbolic groups, then the existence of factor sets in chapter 6: here again, we give different versions of this result (for morphisms to free groups, to torsion-free hyperbolic groups, relatively to a subgroup), which follow from the different versions of the shortening argument.

In the second half, we give the proof of Theorem B. Chapter 7 exposes the result and proves it using proposition C. Finally, the last three chapters are devoted to the proof of proposition C.

Chapter 3

Basic notions

3.1 Actions on simplicial trees and graphs of groups

The notion of graph of groups can be seen as a generalisation of the notion of amalgamated product. Similarly to amalgams, it enables us to understand the structure of a group, and to reduce questions about a group to questions about a finite number of its subgroups. The theory was developed by Jean-Pierre Serre and Hyman Bass, and the main reference is [Ser83]. In this section, we define graphs of groups, and explain the correspondence between them and actions on simplicial trees. We then define a few simple operations that can be applied to a graph of groups, and that will be of use later.

Definition of a graph of groups. We will use the definitions and results of [Ser83] on graphs of groups. Recall that a **graph of groups** Λ is given by

- an underlying oriented graph (that we also denote by Λ), with vertex set $V(\Lambda)$ and edge set $E(\Lambda)$, which is endowed with
 - an involution $\bar{\cdot} : E(\Lambda) \rightarrow E(\Lambda)$ such that $\bar{\bar{y}} \neq y$;
 - applications $o : E(\Lambda) \rightarrow V(\Lambda)$ and $t : E(\Lambda) \rightarrow V(\Lambda)$ which associate to each edge its endpoints, and such that for any $e \in E(\Lambda)$, we have $o(e) = t(\bar{e})$;
- a collection of groups $\{G_v\}_{v \in V(\Lambda)}$, and a collection of groups $\{G_e\}_{e \in E(\Lambda)}$, such that if $e \in E(\Lambda)$, then $G_e = G_{\bar{e}}$;
- injective group morphisms $i_e : G_e \rightarrow G_{t(e)}$ for each e in $E(\Lambda)$.

Pick a maximal subtree Λ_0 in the graph underlying Λ . The **fundamental group** $\pi_1(\Lambda)$ of the graph of groups Λ is defined as the group generated by the groups G_v for $v \in V(\Lambda)$ together with the set $\{t_e \mid e \in E(\Lambda)\}$, with the following relations added:

- $t_e = t_{\bar{e}}^{-1}$ for every edge e of Λ ;
- $t_e = 1$ for every edge e of Λ_0 ;
- $i_e(g) = t_e i_{\bar{e}}(g) t_e^{-1}$, for every edge e of Λ , and every element g of G_e .

Note that the isomorphism class of the fundamental group of Λ does not depend on the choice of the maximal subtree Λ_0 , which justifies the notation $\pi_1(\Lambda)$.

Graphs of groups and actions on trees. Suppose that a group G acts on a simplicial tree T without inversions, so that if an element of G stabilises an edge of T , it fixes it pointwise. Then there is a graph of groups associated to this action, whose fundamental group is isomorphic to G . Such an isomorphism can be built as follows.

We choose a fundamental domain T_0 of the action of G on (the topological realisation) of T , namely a connected subspace of T which contains exactly one point in each orbit. If e is an edge whose interior lies in T_0 note that we must have $o(e) \in T_0$ or $t(e) \in T_0$ by connectedness. A **Bass-Serre element** corresponding to e is an element γ_e of G such that

- $\gamma_e \cdot o(e) \in T_0$ if $o(e) \notin T_0$;
- $\gamma_e^{-1} \cdot t(e) \in T_0$ if $t(e) \notin T_0$;
- 1 otherwise.

Note that the definition implies that for any edge e of T_0 , we have $\gamma_{\bar{e}} = \gamma_e^{-1}$.

Suppose we chose a fundamental domain T_0 , and for each edge e of T_0 , a Bass-Serre element γ_e for the action of G on T . Denote by π the projection $T \rightarrow G \backslash T$. Consider the graph of groups Λ with underlying graph $G \backslash T$, given by

- $G_{\pi(e)} = \text{Stab}_G(e)$ for each edge e whose interior is in T_0 ,
- $G_{\pi(v)} = \text{Stab}_G(v)$ for each vertex v in T_0 ,
- if e is an edge whose interior lies in T_0 , the map $i_{\pi(e)} : G_{\pi(e)} \rightarrow G_{\pi(t(e))}$ is given by $g \mapsto g$ if $t(e) \in T_0$, and by $g \mapsto \gamma_e^{-1} g \gamma_e$ if $t(e) \notin T_0$.

Théorème 13 in [Ser83] tells us that the map from $\pi_1(\Lambda)$ to G defined by $g \mapsto g$ for $g \in G_{\pi(v)}$, and $t_e \mapsto \gamma_e$, is an isomorphism. Note that the choice of a different fundamental domain, or of different Bass-Serre elements, gives us a different isomorphism.

Conversely, the fundamental group G of a graph of groups Λ acts on a tree that we denote T_Λ , the **Bass-Serre tree corresponding to Λ** , in such a way that the graph of groups associated to this action is the original graph of groups Λ . We call **vertex groups** and **edge groups** of Λ all the stabilisers of a vertex, respectively of an edge, of the tree T associated to Λ . In other words, the vertex groups of Λ are the conjugates in G of the groups G_v for v in $V(\Lambda)$, and the edge groups of Λ are the conjugates in G of the groups G_e for e in $E(\Lambda)$.

If Λ is a graph of groups with fundamental group G , we say that Λ is a **splitting** for G . If all the edge groups of Λ are cyclic, or abelian, we say that Λ is a **cyclic splitting**, respectively an **abelian splitting** for G . If the underlying graph of Λ has only one edge, we call Λ a **one edge splitting** for G : it gives for G a structure of amalgamated product (if Λ has two vertices) or of HNN extension (if Λ has only one vertex).

Refining graphs of groups. Suppose we are given a minimal action of a group G on a tree T (i.e. T has no proper G -invariant subtree), with corresponding graph of groups Λ . Let v be a vertex of T , and suppose that the stabiliser G_v of v has a minimal action on a tree T_v , with corresponding graph of groups Γ . Suppose moreover that for any edge e of T adjacent to v , the stabiliser of e fixes a vertex v_e in the action of G_v on T_v .

We build a new G -tree T' from T and T_v . To do this, replace v by T_v , attaching each adjacent edge e to v_e . Extend equivariantly to get T' . There is a canonical action of G on T' that we call the **refinement** of the action of G on T_Λ by the action of G_v on T_Γ . It is minimal, and if the original actions weren't both trivial, it is non-trivial. We call the corresponding graph of group the **refinement** of Λ by Γ .

Quotients of graphs of groups. Let Λ be a graph of group with fundamental group G . Suppose that for each vertex v of Λ with corresponding group G_v , we are given a surjective morphism $q_v : G_v \rightarrow q_v(G_v)$, whose restriction to each $i_e(G_e)$ contained in G_v is injective.

We build a graph of groups Λ_q from Λ as follows: for each vertex v of Λ , we replace the corresponding vertex group G_v by $q_v(G_v)$, and for each edge e such that $t(e) = v$, we replace i_e by $q_v \circ i_e$. The fundamental group of Λ_q is obtained by quotienting G by the smallest normal subgroup of G containing the kernels of all the q_v . We call the graph of groups Λ_q the **quotient of Λ by the maps $(q_v)_{v \in V(\Lambda)}$** .

Extending vertex automorphisms. Let Λ be a graph of groups with fundamental group G . Let ϕ_v be an automorphism of the vertex group G_v corresponding to some vertex v of Λ . Suppose that for every edge e of Λ which is adjacent to v , there exists an element g_e of G_v such that on $i_e(G_e)$, the map ϕ_v restricts to conjugation by g_e . Then we can extend ϕ_v to an automorphism of G . To see this, start by picking a maximal subtree of Λ . For any vertex w of Λ , if the path between w and v in this maximal subtree ends by an edge e , define ϕ to be conjugation by g_e on G_w . Suppose now that t is the generating element corresponding to an edge f which is not in the maximal subtree, with $o(f) = w$ and $t(f) = w'$. Suppose that on G_w and $G_{w'}$ respectively, we defined ϕ as conjugation by some element g_e and $g_{e'}$ respectively. Then we set $\phi(t) = g_e t g_{e'}^{-1}$. It is straightforward to check that ϕ is well defined, and that it is an automorphism. We call this a **standard extension** of ϕ_v to G . Such an element ϕ of $\text{Aut}(G)$ is called a vertex automorphism of G relative to Λ .

3.2 Hyperbolic metric spaces

The notion of hyperbolicity, originally related to curvature of Riemannian manifolds, was extended to general metric spaces by Gromov. We give here several characterisations of hyperbolicity, as well as some basic results. The references here are [GdlH90] and [CDP90].

Definition 3.1: (δ -hyperbolic metric space) *Let (X, d) be a geodesic metric space. A geodesic triangle $\Delta(x, y, z) = [x, y] \cup [y, z] \cup [z, x]$ in X is said to be δ -thin if any point p of $\Delta(x, y, z)$ is in the closed δ -neighbourhood of at least two faces. We say that (X, d) is δ -hyperbolic if all geodesic triangles are δ -thin.*

Given a geodesic triangle $\Delta(x, y, z)$ in a metric space, let $T(x', y', z')$ be the unique tripod whose endpoints x', y', z' are at distances $d(x, y) = d(x', y')$, $d(y, z) = d(y', z')$ and $d(x, z) = d(x', z')$. There exist a unique map $p_\Delta : \Delta(x, y, z) \rightarrow T(x', y', z')$ which restricts to an isometry on each side of the triangle. The following lemma gives us an equivalent definition for δ -hyperbolicity (see proposition 21 of chapter 2 of [GdlH90]).

Lemma 3.2: *The geodesic metric space (X, d) is δ -hyperbolic if and only if for any geodesic triangle $\Delta(x, y, z)$ in X , and for any points z, z' of $\Delta(x, y, z)$ such that $p_\Delta(z) = p_\Delta(z')$, we have $d(z, z') \leq \delta$.*

The following lemma is an easy consequence of this characterisation.

Lemma 3.3: *Let Z be a geodesic quadrilateral contained in X . There exists a simplicial tree $T(Z)$, and a map $p : Z \rightarrow T(Z)$, whose restriction to each side of the quadrilateral is an isometry, and such that for any two points z, z' of Z we have*

$$d(p(z), p(z')) \leq d(z, z') \leq d(p(z), p(z')) + 2\delta.$$

Remark that in the previous lemma, we may assume that the tree $T(Z)$ is spanned by the images u, v, w, x of the vertices of the geodesic quadrilateral. In such a tree, the following inequality

holds:

$$d(u, v) + d(w, x) \leq \max\{d(u, w) + d(v, x); d(u, x) + d(v, w)\}.$$

This suggests yet another characterisation of hyperbolicity, that is also found in [GdlH90]. It is slightly more general in that it makes sense even if the spaces we consider are not geodesic. Note the changes in the hyperbolicity constants.

Lemma 3.4: *If (X, d) is a geodesic metric space such that for any four points v, w, y, z of X we have*

$$d(v, w) + d(y, z) \leq \max\{d(v, y) + d(w, z), d(v, z) + d(w, y)\} + \delta$$

then (X, d) is 2δ -hyperbolic. Conversely, if X is δ -hyperbolic then for any four points v, w, y, z of X we have

$$d(v, w) + d(y, z) \leq \max\{d(v, y) + d(w, z), d(v, z) + d(w, y)\} + 4\delta.$$

We will also need some properties of isometries $f : X \rightarrow X$ of a geodesic δ -hyperbolic space X . The following lemma will prove very useful.

Lemma 3.5: *Let X be a geodesic δ -hyperbolic space X . Let $f : X \rightarrow X$ be an isometry. Suppose x and y are two points of X . Let $t \mapsto v(t)$ be a geodesic parametrisation of a geodesic segment $[x, y]$ for which $v(-T) = x$ and $v(T) = y$.*

Suppose that $d(x, f(x)) + d(y, f(y)) < 2d(x, y) - 4\delta$. Then there exists a real number λ such that $|\lambda| \leq \max\{d(x, f(x)), d(y, f(y))\}$, and for $|t| < T - \max\{d(x, f(x)), d(y, f(y))\}$, we have

$$d(f(v(t)), v(t + \lambda)) < 2\delta \text{ and } d(f^{-1}(v(t)), v(t - \lambda)) < 2\delta.$$

This lemma motivates the following definition

Definition 3.6: (quasitranslation) *Let $t \mapsto x(t)$ be a possibly infinite geodesic arc I in a geodesic δ -hyperbolic space X . We say that the map $f : X \rightarrow X$ acts as an η -translation of length λ on a subarc J of I if for any point $x(t)$ of J , $x(t + \lambda)$ is defined and $d(f(x(t)), x(t + \lambda)) < \eta$.*

Thus lemma 3.5 says that if a map f moves two points of X by a distance which is smaller than the distance between them, it acts as a 2δ -quasitranslation far from the endpoints of a geodesic segment between these two points. Let us prove it.

Proof. Choose geodesic segments $[x, f(x)]$ and $[y, f(y)]$. Consider the geodesic quadrilateral Z formed by these segments together with $[x, y]$ and $f([x, y])$.

We apply lemma 3.3 to get a simplicial tree $T(Z)$, and a map $p : Z \rightarrow T(Z)$. We have

$$\begin{aligned} d(p(x), p(f(x))) + d(p(y), p(f(y))) &\leq d(x, f(x)) + d(y, f(y)) \\ &< 2d(x, y) - 4\delta \\ &= d(x, y) + d(f(x), f(y)) - 4\delta \\ &\leq d(p(x), p(y)) + d(p(f(x)), p(f(y))) \end{aligned}$$

so the segments $[p(x), p(f(x))]$ and $[p(y), p(f(y))]$ do not intersect in $T(Z)$. Denote by m and m' the points for which $[m, m']$ is the shortest path joining these two segments. Let $a = d(p(x), m)$ and $b = d(p(f(x)), m)$, and let also $a' = d(p(y), m')$ and $b' = d(p(f(y)), m')$.

If $t \in [-T, T]$, the point $f(v(t))$ lies on $[f(x), f(y)]$ at a distance $t + T$ of $f(x)$ and $T - t$ of $f(y)$. Thus its image by p lies at a distance $t + T$ of $p(f(x))$, and $T - t$ of $p(f(y))$. If $t \in [-T + b, T - b']$, this implies that it lies in $[m, m']$. In this case, its distance to $p(x)$ is $t - T - b + a$. Let $\lambda = a - b$: the point $p(v(t + \lambda))$ lies in $p([x, y])$ at a distance $t - T - b + a$ of $p(x)$, so it lies on $[m, m']$ and we have $p(f(v(t))) = p(v(t + \lambda))$. We get

$$d(f(v(t)), v(t + \lambda)) < 2\delta.$$

Note now that $d(p(x), p(y)) = a + d(m, m') + a'$, and that $d(p(f(x)), p(f(y))) = b + d(m, m') + b'$ so that $\lambda = a - b = b' - a'$. If $t' \in [-T + a, T - a']$, we let $t = t' - \lambda$: then $t \in [-T + b, T - b']$ so we have

$$d(v(t' - \lambda), f^{-1}(v(t'))) = d(f(v(t' - \lambda)), v(t')) = d(f(v(t)), v(t + \lambda)) < 2\delta.$$

Finally, note that $a + b = d(p(x), p(f(x))) = d(x, f(x))$. Thus a and b are at most $d(x, f(x))$, and a' and b' are at most $d(y, f(y))$. This implies in particular that $|\lambda| = |a - b|$ is bounded by $\max\{d(x, f(x)), d(y, f(y))\}$, and if $|t| < T - \max\{d(x, f(x)), d(y, f(y))\}$ we have $t \in [-T + b, T - b'] \cap [-T + a, T - a']$ so

$$d(f(v(t)), v(t + \lambda)) < 2\delta \text{ and } d(f^{-1}(v(t)), v(t - \lambda)) < 2\delta.$$

□

A geodesic hyperbolic metric space X can be compactified by the addition of a boundary ∂X (see chapter 2 of [CDP90]). The closure of a geodesic of X in $X \cup \partial X$ intersects the boundary in two points, called the points at infinity of this geodesic. We have (this is proposition 2.2 of chapter 2 in [CDP90])

Lemma 3.7: *Two geodesics which have the same points at infinity lie within 8δ of each other.*

Isometries of a hyperbolic metric space can be classified into three types: elliptic, parabolic and hyperbolic. We will only be interested in the latter.

Definition 3.8: (hyperbolic isometry) *An isometry $f : X \rightarrow X$ is hyperbolic if there exists a point x in X for which the map from \mathbb{Z} to X defined by $n \mapsto f^n(x)$ is a quasi-isometry.*

It is easy to see that in this definition we can replace "there exists a point x " by "for any point x ". The quasi-isometry $n \mapsto f^n(x)$ defines two points of the boundary, the limits $\lim_{n \rightarrow \infty} f^n(x)$ and $\lim_{n \rightarrow -\infty} f^n(x)$ that we denote by $f(\infty)$ and $f(-\infty)$ respectively. We know that f can be extended to a homeomorphism of the boundary ∂X : it is clear that this extension fixes these two points. It can be shown that f fixes exactly these two points on the boundary. Moreover, any power of f also fixes $f(-\infty)$ and $f(\infty)$.

Definition 3.9: (axis of a hyperbolic isometry) *We denote by $\text{Ax}(f)$ the union of all the geodesics $t \mapsto x(t)$ of X such that $\lim_{t \rightarrow -\infty} x(t) = f(-\infty)$ and $\lim_{t \rightarrow \infty} x(t) = f(\infty)$.*

Then $\text{Ax}(f)$ is stabilised by f , and if $k \in \mathbb{Z}$, the axis $\text{Ax}(f^k)$ of f^k is just $\text{Ax}(f)$.

Definition 3.10: (translation length) *If $f : X \rightarrow X$ is a hyperbolic isometry, its translation length is $\text{tr}(f) = \inf_{x \in \text{Ax}(f)} d(x, f(x))$.*

Usually the translation length is defined as the infimum of $d(x, f(x))$ for x ranging over the whole space X . However, in the few results we use, this definition is more convenient, and it can be shown that it only differs from the usual translation length by a few δ 's.

Lemma 3.11: *If $x \in \text{Ax}(f)$, we have $d(x, f(x)) \leq \text{tr}(f) + 16\delta$.*

Proof. Suppose x lies on a geodesic L contained in $\text{Ax}(f)$. Let $y \in \text{Ax}(f)$. The map $k \mapsto f^k(y)$ is a quasi-isometry, and its image lies in a 8δ -neighbourhood of L by lemma 3.7. Thus if y_k is a point of L such that $d(y_k, f^k(y)) < 8\delta$, the map $k \mapsto y_k$ is also a quasi-isometry. Thus there is an integer k such that $x \in [y_k, y_{k+1}]$, so that x lies at distance at most 8δ of a geodesic segment $[f^k(y), f^{k+1}(y)]$. Let z be a point of $[f^k(y), f^{k+1}(y)]$ for which $d(x, z) < 8\delta$. Note that

$$d(z, f(z)) \leq d(z, f^{k+1}(y)) + d(f^{k+1}(y), f(z)) = d(f^k(y), z) + d(z, f^{k+1}(y)) = d(f^k(y), f^{k+1}(y))$$

since z lies on a geodesic segment between $f^k(y)$ and $f^{k+1}(y)$. Thus $d(z, f(z)) \leq d(y, f(y))$. Now we have by the triangle inequality

$$d(x, f(x)) \leq 2d(x, z) + d(z, f(z)) \leq 16\delta + d(y, f(y)).$$

This holds for every point y of $\text{Ax}(f)$, so the result holds. \square

Lemma 3.12: *Let $f : X \rightarrow X$ be a hyperbolic isometry. Then f acts as a 20δ -quasitranlation of length $\text{tr}(f)$ on any geodesic contained in $\text{Ax}(f)$.*

Proof. Suppose $t \mapsto v(t)$ parametrises a geodesic in the axis of f , with $\lim_{t \rightarrow -\infty} v(t) = f(-\infty)$, and $\lim_{t \rightarrow \infty} v(t) = f(\infty)$. Let $T \in \mathbb{R}$ such that $4T > 2\text{tr}(f) + 36\delta$.

We then let x and y be the points of $\text{Ax}(f)$ given by $x = v(-T)$ and $y = v(T)$. We have $d(x, f(x)) + d(y, f(y)) \leq 2\text{tr}(f) + 32\delta$ by lemma 3.11, so

$$d(x, f(x)) + d(y, f(y)) < 4T - 4\delta = 2d(x, y) - 4\delta.$$

Therefore we can apply lemma 3.5 to see that for $|t| < T - \text{tr}(f) - 16\delta$, there is a real number λ so that we have

$$d(f(v(t)), v(t + \lambda)) < 2\delta.$$

Moreover, we know that

$$d(v(t), f(v(t))) - 2\delta \leq |\lambda| \leq d(x, f(x))$$

so that $\text{tr}(f) - 2\delta \leq |\lambda| \leq \text{tr}(f) + 16\delta$.

Finally we get $d(f(v(t)), v(t + \text{tr}(f))) < 2\delta + 18\delta \leq 20\delta$. This proves the claim. \square

We show

Lemma 3.13: *Let X be a geodesic δ -hyperbolic space. If f is a hyperbolic isometry $X \rightarrow X$ with $\text{tr}(f) > 12\delta$, for any point x of X , the midpoint m of a geodesic arc $[x, f(x)]$ satisfies $d(m, \text{Ax}(f)) < 4\delta$.*

Proof. Let \bar{x} be a point of $\text{Ax}(f)$ such that $|d(x, \text{Ax}(f)) - d(x, \bar{x})| < \delta$. Note that then,

$$|d(f(x), \text{Ax}(f(x))) - d(f(x), f(\bar{x}))| < \delta.$$

A geodesic segment $[\bar{x}, f(\bar{x})]$ lies within 2δ of $\text{Ax}(f)$.

Consider a geodesic quadrilateral Z formed by $\{x, \bar{x}, f(\bar{x}), f(x)\}$. We apply lemma 3.3 to find a simplicial tree $T(Z)$ and a map $p : Z \rightarrow T(Z)$.

There are only two "combinatorial" possibilities for $p(Z)$. Suppose first that $p([x, f(x)])$ and $p([\bar{x}, f(\bar{x})])$ intersect in a non-trivial segment $[z, z']$. Note then that

$$d(p(x), z) \geq d(x, \text{Ax}(f)) \geq d(x, \bar{x}) - \delta$$

so $d(z, p(x)) \leq 2\delta$. We deduce

$$\begin{aligned} d(z, z') &\geq d(p(\bar{x}), p(f(\bar{x}))) - 2\delta \\ &\geq d(\bar{x}, f(\bar{x})) - 4\delta \\ &\geq \text{tr}(f) - 4\delta \geq 8\delta. \end{aligned}$$

This implies that $p(m) \in [z, z']$. Thus there exists a point y on $[\bar{x}, f(\bar{x})]$ such that $p(y) = p(m)$, and this implies $d(m, \text{Ax}(f)) < 4\delta$.

If $p([x, f(x)])$ and $p([\bar{x}, f(\bar{x})])$ do not intersect, there are points y and y' on $[x, f(x)]$ and $[\bar{x}, f(\bar{x})]$ respectively such that $[p(y), p(y')]$ is the path between them. Then

$$d(p(x), p(\bar{x})) \leq d(x, \bar{x}) \leq d(x, Ax(f)) + \delta \leq d(x, y') + 3\delta \leq d(p(x), p(y')) + 5\delta$$

so that $d(p(\bar{x}), p(y')) < 5\delta$. Similarly, we can see that $d(p(f(\bar{x})), p(y')) < 5\delta$. Thus

$$d(\bar{x}, f(\bar{x})) < d(p(\bar{x}), p(f(\bar{x}))) + 2\delta < d(p(\bar{x}), y') + d(y', p(f(\bar{x}))) + 2\delta < 12\delta.$$

□

We will also need

Lemma 3.14: *Let Γ be a torsion-free hyperbolic group. Denote by X its Cayley graph with respect to some generating set Σ . For $R > 0$, there exists a constant M_R such that for any non-trivial element g , the translation length of g^{M_R} is at least R .*

The proof follows that of proposition 3.1 in [Del96].

Proof. Denote by δ a hyperbolicity constant for X , we can assume without loss of generality that it is an integer. We fix an order on the generating set Σ , and then order words in the elements of Σ lexicographically. We say that a geodesic L in X is special if for any two points g and g' on L , the word in Σ corresponding to the segment of L between g and g' is minimal among words representing $g^{-1}g'$. It can be shown by following the proof of proposition 2.2 in chapter 2 of [CDP90] that any two points on the boundary of X are joined by a least one special geodesic. If we pick two points in the boundary of X , and two disjoint balls of radius 8δ centred on a special geodesic joining them, we see by lemma 3.7 that any other special geodesic must pass through both these balls. On the other hand, given any pair of points x, y in X , there is at most one special geodesic containing both x and y . Thus the number of special geodesics joining any two points on the boundary is bounded by $|B_{8\delta}(X)|^2$.

Let now g be a non-trivial element of G . It is a known result (see Théorème 3.3 and Théorème 3.4 in chapter 9 of [CDP90]) that any non-trivial element of a hyperbolic group acts hyperbolically on its Cayley graph. Thus g fixes two points on the boundary. The image by g of a special geodesic joining them must also be a special geodesic: g permutes the set of special geodesics. Thus if $M = |B_{8\delta}(X)|^2$, we know that g^M fixes all the special geodesics. In particular, its restriction to a special geodesic is a translation, of length at least 1 since the distance function has integer values. Let $M_R = M(R + 16\delta)$: the element g^{M_R} restricted to a special geodesic is a translation of length at least $R + 16\delta$. We conclude by applying lemma 3.11. □

3.3 Limits of metric spaces

In all this section, G is a group endowed with a finite generating set Σ_G . We want to define a topology on a set of pointed metric G -spaces, and to give a criterion for a sequence in such a set to admit a convergent subsequence. Then, we look at the particular case where the metric spaces in the sequence are all hyperbolic, and we see under which conditions the limit is a real tree.

3.3.1 The Gromov-Hausdorff topology

Let $\mathcal{A}(G)$ be a set of pointed metric G -spaces, that is, metric spaces endowed with an action of G by isometries. We want to define a topology on $\mathcal{A}(G)$ called the equivariant Gromov-Hausdorff topology. It is a generalisation of the Gromov-Hausdorff topology on a set of compact metric spaces (see [Pau88]).

Definition 3.15: (ϵ -approximation) Let $(K, d), (K', d')$ be two compact metric spaces. Let $\epsilon < 0$. An ϵ -approximation between K and K' is a binary relation $R \subset K \times K'$ whose projections on K and K' are surjective, and such that for $x, y \in K$ and $x', y' \in K'$, if xRx' and yRy' , then

$$|d(x, y) - d'(x', y')| < \epsilon.$$

The Gromov-Hausdorff distance between two compact metric spaces is the infimum of the set of ϵ for which an ϵ -approximation exists. This can be generalised to a topology on sets of non-compact pointed metric spaces, where a sequence (X_n, x_n) converges to (X, x) if the Gromov Hausdorff distance between the ball centred on x_n of radius n and the ball of radius n centred on x tends to 0 as n tends to infinity. We want to further generalise this to metric spaces endowed with an action of G .

Definition 3.16: (neighbourhoods $N(G_0, K, \epsilon)$) Let (X, x) be an element of $\mathcal{A}(G)$. Given a finite subset K of X , a finite subset G_0 of G , and $\epsilon > 0$, we say that an element (X', x') of $\mathcal{A}(G)$ is in $N(G_0, K, \epsilon)(X, x)$ if there exists a finite subset K' of X' , and an ϵ -approximation R between K and K' such that if $y, z \in K$ and $y', z' \in K'$, and if yRy', zRz' , then for any element g of G_0

$$|d(y, g \cdot z) - d'(y', g \cdot z')| < \epsilon$$

Definition 3.17: (equivariant Gromov-Hausdorff topology) The equivariant Gromov-Hausdorff topology on $\mathcal{A}(G)$ is the topology generated by the neighbourhoods of the form $N(G_0, K, \epsilon)$ for $(X, x) \in \mathcal{A}(G)$, for G_0 a finite subset of G , for K a finite subset of X , and $\epsilon > 0$.

3.3.2 Ultraproducts and limit of sequences

We want to give a sufficient condition on a sequence (X_n, x_n) in $\mathcal{A}(G)$ to ensure that it contains a convergent subsequence. To build a limit, we will need the following tools.

Definition 3.18: (filter) A filter F on \mathbb{N} is a non-empty subset of $\mathcal{P}(\mathbb{N})$ such that

- if $A, B \in F$, then $A \cap B \in F$;
- if $A \in F$ and if $A \subseteq B$, then $B \in F$.

Example 3.19:

- the principal filter over an element n of \mathbb{N} is the set of all subsets of \mathbb{N} which contain n ;
- the Frechet filter is the set of all cofinite subset of \mathbb{N} .

Definition 3.20: (ultrafilter) A filter is an ultrafilter if it is maximal for inclusion.

Remark 3.21: A filter is an ultrafilter if and only if for each subset A of \mathbb{N} , it contains exactly one of A , $\mathbb{N} - A$.

A principal filter is an ultrafilter. The Frechet filter isn't an ultrafilter (it contains neither the set of even numbers, nor the set of all odd numbers its complement).

By applying Zorn's lemma, we can show that any filter is contained in an ultrafilter. Thus we can enlarge the Frechet filter to an ultrafilter which is easily seen to be non-principal. Conversely, any non-principal ultrafilter contains the Frechet filter.

Definition 3.22: (limits with respect to ω) Let ω be a non-principal ultrafilter. We say that the sequence $(u_n)_{n \in \mathbb{N}}$ of real numbers tends to u with respect to ω if for all $\epsilon > 0$, the set $\{n \in \mathbb{N} \mid |u_n - u| < \epsilon\}$ is in ω . We denote this $\lim_{\omega} u_n = u$.

If a sequence of reals $(u_n)_{n \in \mathbb{N}}$ tends to a limit u as n tends to infinity, then it tends to u with respect to any non-principal ultrafilter ω . If a sequence of reals $(u_n)_{n \in \mathbb{N}}$ is bounded, then it admits a limit with respect to any non-principal ultrafilter ω . In fact, if a sequence $(u_n)_{n \in \mathbb{N}}$ is bounded with respect to ω , namely if there exists a constant M such that $\{n \in \mathbb{N} \mid u_n < M\}$ lies in ω , then it tends to a limit with respect to ω . If this is not the case, we say that u_n tends to ∞ with respect to ω .

Let $(X_n, x_n)_{n \in \mathbb{N}}$ be a sequence of pointed metric spaces. Let ω be a non-principal ultrafilter. Define X_∞ to be the set

$$\{(y_n)_{n \in \mathbb{N}} \in \prod_{n \in \mathbb{N}} X_n \mid \lim_{\omega} d_n(x_n, y_n) < \infty\}$$

quotiented by the following equivalence relation:

$$(y_n)_{n \in \mathbb{N}} \sim (z_n)_{n \in \mathbb{N}} \iff \{n \in \mathbb{N} \mid y_n = z_n\} \in \omega.$$

We will usually denote (y_n) both the sequence and its class for \sim . We choose (x_n) as a basepoint for X_∞ .

Now for all $(y_n), (z_n) \in X_\infty$ we set:

$$d_\infty((y_n), (z_n)) = \lim_{\omega} d_n(y_n, z_n).$$

This limit exists since $d_n(y_n, z_n) \leq d_n(y_n, x_n) + d_n(x_n, z_n)$, and both $d_n(y_n, x_n)$ and $d_n(x_n, z_n)$ are bounded with respect to ω by hypothesis. This defines a pseudometric on X_∞ . Denote by (X_ω, x_ω) the quotient of X_∞ by the equivalence relation given by $d_\infty((y_n), (z_n)) = 0$, and by d_ω the metric on X_ω . Again we abuse notations and denote the equivalence class of a point (y_n) of X_∞ by (y_n) .

Suppose now that each X_n is endowed with an action of G such that, for any element a of the generating set Σ_G , the sequence $(d_n(x_n, a \cdot x_n))_{n \in \mathbb{N}}$ is bounded with respect to ω . Then for any element (y_n) of X_ω , and any element $g = a_1 \dots a_s$ with $a_i \in \Sigma_G$, we have

$$\begin{aligned} d_n(x_n, g \cdot y_n) &\leq d_n(x_n, g \cdot x_n) + d_n(g \cdot x_n, g \cdot y_n) \\ &\leq \sum_{i=1}^s d_n(x_n, a_i \cdot x_n) + d_n(x_n, y_n) \end{aligned}$$

Thus the sequence $d_n(x_n, g \cdot y_n)$ is bounded with respect to ω , so the sequence $(g \cdot y_n)_{n \in \mathbb{N}}$ defines a point in X_ω . It is straightforward to check that $(g, (y_n)_n) \mapsto (g \cdot y_n)$ gives an action of the group G on X_ω by isometries.

Definition 3.23: (ultraproduct of a sequence of G -spaces) *Let (X_n, x_n) be a sequence in $\mathcal{A}(G)$ for which there exists a non-principal ultrafilter ω , such that for any element g of the generating set Σ_G , the sequence $(d_n(x_n, g \cdot x_n))_{n \in \mathbb{N}}$ is bounded with respect to a non-principal ultrafilter ω . The space (X_ω, x_ω) endowed with the action of G given by*

$$(g, (y_n)_n) \mapsto (g \cdot y_n)$$

is called the ultraproduct of the G -spaces (X_n, x_n) with respect to ω .

Ultraproducts are natural limits for sequences of G -spaces.

Lemma 3.24: *Let $(X_n, x_n)_{n \in \mathbb{N}}$ be a sequence in $\mathcal{A}(G)$. Let ω be a non-principal ultrafilter such that (X_ω, x_ω) is defined and lies in $\mathcal{A}(G)$. Then (X_ω, x_ω) lies in the closure of $\{(X_n, x_n) \mid n \in \mathbb{N}\}$ in $\mathcal{A}(G)$ with respect to the equivariant Gromov-Hausdorff topology.*

Proof. Let g be an element of G , let $y = (y_n)$ and $z = (z_n)$ be points in X_ω , and let $\epsilon > 0$. We get from the definition of d_ω that the set $A^\epsilon(g, y, z)$ defined by

$$A^\epsilon(g, y, z) = \{n \in \mathbb{N} \mid |d_n(y_n, g \cdot z_n) - d_\omega(y, g \cdot z)| < \epsilon\}$$

lies in ω .

For any finite subset G_0 of G , any finite subset K of X_ω and any $\epsilon > 0$, the intersection $A^\epsilon(G_0, K)$ of all the sets of the form $A^\epsilon(g, y, z)$ for g in G_0 and y, z in K is still in ω , in particular it is not empty. But (X_n, x_n) lies in $N(G_0, K, \epsilon)$ precisely if n lies in $A^\epsilon(g, y, z)$ for all g in G_0 and y, z in K . This proves the lemma. \square

Note that this lemma implies in particular that some subsequence of $(X_n, x_n)_{n \in \mathbb{N}}$ tends to (X_ω, x_ω) . Thus, for a sequence $(X_n, x_n)_{n \in \mathbb{N}}$ to admit a convergent subsequence, it is enough that there exist a non-principal ultrafilter ω with respect to which the sequence $d_n(x_n, g \cdot x_n)$ is bounded for every element g of Σ_G . This boundedness condition might not be satisfied, but by rescaling properly we can overcome this problem. For this, define

Definition 3.25: (length of an action) *If (X, x) is a pointed G -space, we define the length of the action to be*

$$l(X, x) = \max_{g \in \Sigma_G} d(x, g \cdot x),$$

the maximal displacement of the basepoint by a generator.

If X is a space endowed with a metric d , and a a positive real, we denote by $\frac{1}{a}X$ the metric space whose underlying set is X , and whose metric is d/a .

Remark 3.26: *Let $(X_n, x_n)_{n \in \mathbb{N}}$ be a sequence of G -spaces: if we rescale the metric on X_n by $l_n = l(X_n, x_n)$, the sequence $(\frac{1}{l_n}(X_n, x_n))_{n \in \mathbb{N}}$ satisfies the condition of boundedness which ensures that the ultraproduct of the spaces $\frac{1}{l_n}(X_n, x_n)$ is defined with respect to any non-principal ultrafilter ω . Thus, up to rescaling by the action lengths, any sequence of metric G -spaces admits a convergent subsequence. This trick will prove very useful in the sequel.*

3.3.3 Limits of pointed hyperbolic G -spaces

We will be interested in limits of sequences of hyperbolic G -spaces. We get the following result about the ultraproduct of path-connected hyperbolic spaces:

Lemma 3.27: *Let ω be a non-principal ultrafilter. Let $(X_n, x_n)_{n \in \mathbb{N}}$ be a sequence in $\mathcal{A}(G)$. If each X_n is a geodesic δ_n -hyperbolic space, if $\lim_\omega \delta_n = 0$, and if (X_ω, x_ω) is defined, it is a real G -tree.*

Proof. To see that X_ω is a real tree, it is enough to see that it is 0-hyperbolic and path connected. For a proof of this, see for example Théorème 4.1 in chapter 3 of [CDP90].

Recall the characterisation of hyperbolicity given by lemma 3.4. If $(v_n), (w_n), (y_n), (z_n) \in X_\omega$, we have:

$$\begin{aligned} d_\omega((v_n), (w_n)) + d_\omega((y_n), (z_n)) &= \lim_\omega (d_n(v_n, w_n) + d_n(y_n, z_n)) \\ &\leq \lim_\omega (\max\{d_n(v_n, y_n) + d_n(w_n, z_n), d_n(v_n, z_n) + d_n(y_n, w_n)\} + 2\delta_n) \\ &\leq \max\{\lim_\omega [d_n(v_n, y_n) + d_n(w_n, z_n)], \lim_\omega [d_n(v_n, z_n) + d_n(y_n, w_n)]\} + 2 \lim_\omega \delta_n \\ &= \max\{d_\omega((v_n), (y_n)) + d_\omega((w_n), (z_n)), d_\omega((v_n), (z_n)) + d_\omega((y_n), (w_n))\} \end{aligned}$$

which proves X_ω is 0-hyperbolic.

Let us see that X_ω is path connected. Let $(v_n) \in X_\omega$. We know that $d_n(x_n, v_n)$ is bounded with respect to ω , in particular there exists a positive constant M such that the set $A = \{n \mid d_n(x_n, v_n) < M\}$ lies in ω . If $n \in A$, let $t \mapsto v_n(t)$ for $t \in [0, M]$ be a 1-Lipschitz path from v_n to x_n (it exists since X_n is a geodesic metric space). If $n \notin A$, let $v_n(t)$ be the constant path. Then the map $t \mapsto (v_n(t))$ is continuous since the v_n are all 1-Lipschitz, and it is a path from (v_n) to (x_n) in X_ω . \square

Chapter 4

Shortening argument

The shortening argument has many variants, of which we will present two. The classical result asserts that, given a sequence of morphisms from a freely indecomposable group G to a torsion-free hyperbolic group Γ , either we can 'shorten' some of the morphisms in the sequence, or the stable kernel is non-trivial (see Theorem 4.25). The length of a morphism $f : G \rightarrow \Gamma$ depends on the choice of generating sets for G and Γ , and of a basepoint in the corresponding Cayley graph of Γ . It is the maximal displacement of this basepoint by the image of one of the generating elements chosen for G . Then 'shortening' a morphism $f : G \rightarrow \Gamma$ is just precomposing it by an automorphism σ of G , in such a way that the length of $f \circ \sigma$ is strictly smaller than that of f .

We will also give a relative version of the shortening argument, in which the group G is only assumed to be freely indecomposable relative to a subgroup H , but the morphisms in the sequence are assumed to fix H in the limit (see Theorem 4.33). We call this type of results morphisms shortening results.

In the proof of both the standard and the relative versions of the shortening argument for morphisms, the first step is to construct from the given sequence of maps $G \rightarrow \Gamma$ a sequence of actions on δ_n -hyperbolic spaces X_n .

Definition 4.1: (action $X[h]$ induced by a morphism) *Let G and Γ be groups endowed with finite generating sets Σ_G and Σ_Γ . If $h : G \rightarrow \Gamma$ is a morphism, G acts on Γ by $(g, \gamma) \mapsto h(g)\gamma$. This induces an action of G on the Cayley graph X of Γ with respect to Σ_Γ , giving it a structure of G -space that we denote by $X[h]$.*

We will see that given a sequence of morphisms $h_n : G \rightarrow \Gamma$, by suitably rescaling the metric on $X[h_n]$ and choosing the right basepoint, we get a sequence of actions on hyperbolic spaces which converges to a non-trivial action on a real tree T .

The second step is to prove an "action shortening result" (theorem 4.20 and theorem 4.28 respectively), which tells us that if a sequence of actions on hyperbolic spaces converges to an action on real tree satisfying certain conditions, we can shorten all but finitely many of the actions. This action shortening result should be considered as the heart of the shortening argument; indeed it can be used to prove results about sequences of action which do not necessarily come from morphisms to a free or a hyperbolic group (see for example [Sel97a]).

In this chapter, we prove the "morphism shortening results" using the "action shortening results". In the first section, we explain how to get a sequence of actions converging to a real tree from a sequence of morphisms $G \rightarrow \Gamma$. In the second and the third section, we state the action shortening theorems and use them to prove the morphism shortening results in the classical and the relative case respectively. The fourth section gives a result which is a straightforward consequence of the shortening argument, and which we will use in the proof of the main result of this thesis. The proof of the action shortening result is given in the next chapter.

4.1 Limit of morphisms to a hyperbolic group

For this whole section, let G be a group endowed with a finite generating set Σ_G , and let Γ be a torsion-free hyperbolic group endowed with a finite generating set Σ_Γ . Denote by X the Cayley graph of Γ with respect to Σ_Γ .

This lemma gives the setting which we will consider in this section. Recall that in definition 3.25, the length of an action of G on a pointed metric space was defined as the minimal displacement of the basepoint by an element of Σ_G .

Lemma 4.2: *Let $h_n : G \rightarrow \Gamma$ be a sequence of morphisms. Suppose that there is a sequence of basepoints x_n for the spaces $X[h_n]$, and a non-principal ultrafilter ω , for which the sequence $l_n = l(X[h_n], x_n)$ tends to infinity with respect to ω . Then the ultraproduct $(X[h, \omega], x_\omega)$ of the spaces $\frac{1}{l_n}(X[h_n], x_n)$ is a real G -tree.*

Proof. By remark 3.26 the ultraproduct of the spaces $\frac{1}{l_n}(X[h_n], x_n)$ is defined. Moreover, the space $X[h_n]$ is δ -hyperbolic, so that $\frac{1}{l_n}(X[h_n], x_n)$ is δ/l_n -hyperbolic. Now δ/l_n tends to 0 with respect to ω , so by lemma 3.27, the ultraproduct $(X[h, \omega], x_\omega)$ of the spaces $\frac{1}{l_n}(X[h_n], x_n)$ is defined, and is a real G -tree. \square

We want to study some properties of such a limit action. The proofs we give are inspired by those found in [Pau97]. The following lemma will prove useful.

Lemma 4.3: *Suppose we are in the setting described in 4.2. For any pair of points y, z of $(X[h, \omega], x_\omega)$, we define D_{yz} as the set*

$$\{aba^{-1}b^{-1} \mid a, b \in G \text{ such that } d(y, a \cdot y), d(y, b \cdot y), d(z, a \cdot z), d(z, b \cdot z) < d(y, z)/12\}.$$

Let R be the cardinal of the ball of radius 8δ in Γ (where Γ is endowed with the word metric associated to Σ_Γ). If y, z are points in $(X[h, \omega], x_\omega)$, and if D_{yz}^0 is a finite subset of D_{yz} , the cardinal of $h_n(D_{yz}^0)$ is bounded by R with respect to ω .

Proof. Let y, z be points of X . Let A_{yz}^0 be a finite set of pairs (a, b) of elements of G such that $d(y, a \cdot y), d(y, b \cdot y), d(z, a \cdot z), d(z, b \cdot z) < d(y, z)/12$ and $[a, b] \in D_{yz}^0$. Let $\epsilon < d(y, z)/20$. There is a set U in ω such that for any n in U , there is an ϵ approximation between T and X_n relative to y, z , and to the elements of the pairs which lie in A_{yz}^0 . We can assume moreover that if $n \in U$ and such that $2\delta_n < \epsilon$. Fix an index n in U , and let y_n and z_n be some points approximating y and z .

Note that $d(y, z) < 10d_n(y_n, z_n)/9$ so $2\delta_n < \epsilon < d_n(y_n, z_n)/9$. Let $(a, b) \in A_{yz}^0$. Suppose one of $h_n(a)$ or $h_n(b)$ is trivial: then $h_n(aba^{-1}b^{-1}) = 1$. Since we want to bound the cardinal of $\{h_n([a, b]) \mid (a, b) \in A_{yz}^0\}$, we can ignore this case.

Consequently, the elements $h_n(a)$ and $h_n(b)$ act hyperbolically on X . Let $[y_n, z_n]$ be a geodesic segment, and let $t \mapsto w(t)$ be a geodesic parametrisation $[-T, T] \rightarrow [y_n, z_n]$. The elements $h_n(a)$ and $h_n(b)$ move y_n and z_n by a distance which is small compared to the distance between them. More precisely, we have

$$d_n(y_n, a \cdot y_n) < d(y, a \cdot y) + \epsilon < d(y, z)/10 + \epsilon < 3d(y, z)/20 < d_n(y_n, z_n)/6,$$

so in particular $d_n(y_n, a \cdot y_n) < d_n(y_n, z_n) - 2\delta_n$. Similarly we show that $d_n(y_n, b \cdot y_n)$, $d_n(z_n, a \cdot z_n)$ and $d_n(z_n, b \cdot z_n)$ are all smaller than $3d(y, z)/20$, so that they are smaller than $d_n(y_n, z_n) - 2\delta_n$, and lemma 3.5 can be applied to the isometries given by a and b on X_n .

Thus, there exists reals λ_a and λ_b , with $|\lambda_a| < \max\{d_n(y_n, a \cdot y_n), d_n(z_n, a \cdot z_n)\} < 3d(y, z)/20 < T/3$ and $|\lambda_b| < T/3$ such that a, b, a^{-1} and b^{-1} act as $2\delta_n$ -quasitranslations of length $\lambda_a, \lambda_b, -\lambda_a$ and $-\lambda_b$ on the subsegment $\{w(t) \mid |t| < 2T/3\}$ of $[z_n, y_n]$.

The idea is now that, up to a few δ_n 's, a and b commute on a subsegment of $[y_n, z_n]$, so their commutator does not move the midpoint $w(0)$ of $[y_n, z_n]$ by more than R .

Since both $|\lambda_a|$ and $|\lambda_b|$ are less than $T/3$, we can apply the inequality given in 3.5 to $t = -\lambda_b$, $t = -\lambda_b - \lambda_a$, and $t = -\lambda_a$. We deduce

$$\begin{aligned} d_n(w(0), aba^{-1}b^{-1} \cdot w(0)) &\leq d_n(w(0), aba^{-1} \cdot w(-\lambda_b)) + 2\delta_n \\ &\leq d_n(w(0), ab \cdot w(-\lambda_a - \lambda_b)) + 4\delta_n \\ &\leq d_n(w(0), a \cdot w(\lambda_b - \lambda_b - \lambda_a)) + 6\delta_n \\ &\leq d_n(w(0), w(\lambda_a + \lambda_b - \lambda_a - \lambda_b)) + 8\delta_n \\ &= d_n(w(0), w(0)) + 8\delta_n = 8\delta_n \end{aligned}$$

This shows that for any n in U , the translates of $w(0)$ under elements of $h_n(D_{yz}^0)$ all lie in the ball of radius $8\delta_n$. Thus, there is a point w of X whose translates by the elements of $h_n(D_{yz}^0)$ lie in a ball of radius 8δ around w . But the action of Γ on its Cayley graph X is free and discrete, and the cardinal of the set of elements translating a point of X by less than a constant C is bounded by the cardinal of the ball of radius C in Γ endowed with the word metric. This proves the claim. \square

Definition 4.4: (stable kernel with respect to an ultrafilter) *The stable kernel with respect to ω of a sequence of morphisms $h_n : G \rightarrow \Gamma$ is the set of elements g of G such that $\{n \mid h_n(g) = 1\}$ lies in ω . We denote it by $\text{Ker}_\omega(h_n)$.*

Suppose we are in the setting of lemma 4.2. Then the stable kernel with respect to ω acts trivially on $(X[h, \omega], x_\omega)$. Indeed, denote by d_n the metric on $\frac{1}{l_n}(X[h_n], x_n)$: we have $d_n = d_{\Sigma_\Gamma}/l_n$. For an element g in the stable kernel of $(h_n)_{n \in \mathbb{N}}$ with respect to ω , and for any point (y_n) in $(X[h, \omega], x_\omega)$, the set of indices n for which the distance $d_n(y_n, g \cdot y_n)$ is zero lies in ω . Thus we have

$$d_\omega((y_n), g \cdot (y_n)) = \lim_{\omega} d_n(y_n, g \cdot y_n) = 0.$$

We now show

Lemma 4.5: *Suppose we are in the setting given by 4.2. The elements of G which stabilise a non-trivial tripod of $(X[h, \omega], x_\omega)$ lie in the stable kernel of $(h_n)_{n \in \mathbb{N}}$ with respect to ω .*

Proof. Let a, b, c be points of $X[h, \omega]$ which form a non-trivial tripod of centre e . Let g be an element of G which fixes this tripod pointwise. Let $\epsilon < \min\{d(a, e), d(b, e), d(c, e)\}/10$. There is an element U_g of ω such that if $n \in U_g$, there is an ϵ -approximation between $X[h, \omega]$ and X_n with respect to a, b, c and g .

Let n be in U_g . Denote by a_n, b_n, c_n the points approximating a, b, c , and let $\Delta(a_n b_n c_n)$ be a geodesic triangle with vertices a_n, b_n, c_n in X_n .

Recall there is a unique map $p_\Delta : \Delta(a_n b_n c_n) \rightarrow Y$, where Y is the unique tripod whose sides have the same lengths as the sides of $\Delta(a_n b_n c_n)$, and the restriction of p_Δ to each face of $\Delta(a_n b_n c_n)$ is an isometry. Let x_n, y_n , and z_n be the points of $[a_n, b_n]$, $[b_n, c_n]$ and $[a_n, c_n]$ respectively such that $p_\Delta(x_n), p_\Delta(y_n)$ and $p_\Delta(z_n)$ are all equal to the centre of the tripod Y . By lemma 3.2, the diameter of $\{x_n, y_n, z_n\}$ is less than δ_n .

We have $d_n(a_n, g \cdot a_n) \leq \epsilon$ and $d_n(b_n, g \cdot b_n) \leq \epsilon$ so far from its endpoints, the path $g \cdot [a_n, b_n]$ by g lies in a $2\delta_n$ -neighbourhood of $[a_n, b_n]$. In particular, we get

$$d_n(g \cdot x_n, [a_n, b_n]) < 2\delta_n$$

and we can find a point x'_n on $[a_n, b_n]$ such that $d(x'_n, g \cdot x_n) < 2\delta_n$. Similarly, we can find points y'_n in $[b_n, c_n]$ for which $d_n(y'_n, g \cdot y_n) < 2\delta_n$ and z'_n in $[a_n, c_n]$ such that $d_n(z'_n, g \cdot z_n) < 2\delta_n$. The

diameter of $\{g \cdot x_n, g \cdot y_n, g \cdot z_n\}$ is less than δ_n so the diameter of $\{x'_n, y'_n, z'_n\}$ is at most $5\delta_n$. Since x'_n, y'_n and z'_n lie on the three different faces of the triangle $\Delta(a_n b_n c_n)$, their images by p_Δ cannot all lie on the same leg of the tripod Y . Thus they are $5\delta_n$ -close to the centre of Y , so that $d(p_\Delta(x'_n), p_\Delta(x_n)) < 5\delta_n$ and we get

$$d_n(g \cdot x_n, x_n) < d_n(x_n, x'_n) + d_n(x'_n, g \cdot x_n) < d(p_\Delta(x'_n), p_\Delta(x_n)) + \delta_n + 2\delta_n < 8\delta_n.$$

Thus any element of G fixing the tripod a, b, c translates x_n by at most $8\delta_n$. But if g fixes the tripod, all its powers do, therefore, the cardinal of the set $\{h_n(g^k)\}_{k \in \mathbb{Z}}$ is bounded by the cardinal of the ball of radius 8δ in Γ endowed with the word metric associated to Σ_Γ . Since it is a subgroup of Γ , and since Γ is torsion-free, it must be trivial. Thus for every n in U_g , we have $h_n(g) = 1$, and g lies in the stable kernel of $(h_n)_{n \in \mathbb{N}}$ with respect to ω . \square

Recall that the elements of the stable kernel of $(h_n)_{n \in \mathbb{N}}$ with respect to ω act trivially on $X[h, \omega]$. The following lemma gives a partial converse.

Lemma 4.6: *Suppose we are in the setting described in 4.2.*

If $(X[h, \omega], x_\omega)$ isn't a line, the kernel of the action of G on $(X[h, \omega], x_\omega)$ is precisely the stable kernel of the sequence $(h_n)_{n \in \mathbb{N}}$ with respect to ω .

If the stable kernel of the sequence $(h_n)_{n \in \mathbb{N}}$ with respect to ω is trivial, and if $(X[h, \omega], x_\omega)$ is a line, then $h_n(G)$ is cyclic for all n .

Proof. Suppose that $X[h, \omega]$ is not a line: it contains a non-trivial tripod, which is fixed by any element which lies in the kernel of the action. But by lemma 4.5, elements of G fixing a tripod in $X[h, \omega]$ lie in the stable kernel of $(h_n)_{n \in \mathbb{N}}$ with respect to ω . The other inclusion has already been proved.

Suppose now $X[h, \omega]$ is a line L , and that $\underline{\text{Ker}}_\omega(h_n) = 1$. Let D_0 be a finite set of commutators of G , and let G_0 be a finite subset of G such that any element in D_0 is a commutator of two elements in G_0 . The elements of G_0 either fix a point of L , or they act by translation: denote by M the maximum of their translation lengths. Let y, z be two points of L such that $d(y, z) > 12M$. Note that D_0 is a subset of the set D_{yz} defined in lemma 4.3, so there is a set U in ω such that for any index n in U , the cardinal of $h_n(D_0)$ is bounded by the constant $R(\Gamma)$. However, there exists also a set U' in ω such that for any index n in U' , the map h_n is injective on D_0 . Since $U \cap U'$ is not empty, we see that the cardinal of D_0 is bounded by $R(\Gamma)$. This shows that the set of commutators of G is finite, so by lemma 1.A in [Pau97], G is virtually abelian. For any index n , the image $h_n(G)$ is virtually abelian. Since it is a subgroup of a torsion-free hyperbolic group, it is in fact cyclic. Indeed, abelian groups in torsion-free hyperbolic groups are infinite cyclic, and virtually cyclic torsion-free groups are cyclic (to see this, show first that the centre of a virtually cyclic group must have finite index, then show that if the centre of a group has finite index, then the derived subgroup must be finite). \square

Finally, we have

Lemma 4.7: *Suppose we are in the setting given by 4.2. Suppose moreover that any virtually abelian subgroup of G is abelian, and that the stable kernel of $(h_n)_{n \in \mathbb{N}}$ with respect to ω is trivial. Then the pointwise stabiliser of an arc of $(X[h, \omega], x_\omega)$, is abelian.*

Proof. Let $[y, z]$ be an arc in $X[h, \omega]$. Let G_1 be a finitely generated subgroup of $\text{Stab}([y, z])$, and suppose that the set D_1 of its commutators is infinite. Let D_0 be a finite subset of D_1 with $|D_0| > R(\Gamma)$.

Note that D_0 lies in D_{yz} , so by lemma 4.3, there is a set U of ω such that for any index n in U , the set $h_n(D_0)$ has cardinal bounded by $R(\Gamma)$. The stable kernel of the sequence $(h_n)_{n \in \mathbb{N}}$ with respect to ω is trivial, so there exists a set U' in ω such that for $n \in U'$, the map h_n is injective on D_0 . Since $U \cap U'$ is not empty, we get a contradiction.

Thus G_1 has a finite set of commutators, so by lemma 1.A in [Pau97], G_1 is virtually abelian. By hypothesis it is in fact abelian. Any finitely generated subgroup of $\text{Stab}([z, y])$ is abelian: it must itself be abelian. \square

4.2 Shortening morphisms in the classical case

4.2.1 Modular group

We start by defining a subgroup of the group of automorphisms $\text{Aut}(G)$ of G called the modular group. We need the following definitions.

Definition 4.8: (Dehn twist) *Let G be a finitely generated group. Suppose Λ is a one edge splitting for G , with edge group C , and let c be an element in the centre of C . A Dehn twist about c is an automorphism ϕ of G defined as follows:*

1. *If $G = A *_C B$, ϕ is the unique automorphism of G which is the identity on A , and conjugation by c on B .*
2. *If $G = A *_C$, and t is a stable letter of this HNN extension, ϕ is the unique automorphism of G which is the identity on A , and sends t to tc .*

If Λ is a graph of groups decomposition of G , the Dehn twists of Λ are the Dehn twists associated to one-edge splittings of G obtained from Λ by collapsing all its edges except one.

Definition 4.9: (generalised Dehn twist) *Suppose G has a graph of groups decomposition Λ , and let A be an abelian vertex group in this decomposition. Let A_1 be the subgroup of A generated by all the incident edge groups. Any automorphism of A which fixes A_1 pointwise can be extended to an automorphism of the whole group, which we call a generalised Dehn twist.*

To define yet another type of automorphisms, we need

Definition 4.10: (maximal boundary subgroups, boundary subgroups, boundary elements) *Let Σ be a surface with boundary. Denote by S its fundamental group (S is a free group).*

To each boundary component of Σ corresponds a conjugacy class of maximal cyclic subgroups of S : we call such groups maximal boundary subgroups. We will refer to generators of maximal boundary subgroups as maximal boundary elements. A non-trivial non-trivial subgroup of a maximal boundary subgroup is a boundary subgroup, and non-trivial elements of such subgroups are boundary elements.

Remark 4.11: *The set of conjugacy classes of the maximal boundary subgroups is in bijection with the set of connected components of the boundary of Σ .*

Definition 4.12: (graph of groups with surfaces) *A graph of groups with surfaces is a graph of groups Γ together with a subset V_S of $V(\Gamma)$ such that any vertex v in V_S satisfies:*

- *there are no loops at v , i.e. no edges both of whose endpoints are v ;*
- *there exists a compact connected surface with boundary Σ which is not a disk, a Möbius band or a cylinder, and such that the vertex group G_v is $S = \pi_1(\Sigma)$;*
- *for each edge e such that $t(e) = v$, the injection $i_e : G_e \hookrightarrow G_v$ maps G_e onto a maximal boundary subgroup of S ;*
- *this induces a bijection between the set of edges $t^{-1}(v)$ and the set of conjugacy classes in S of maximal boundary subgroups of S ;*

The vertices of V_S are called *surface type vertices*. A vertex of the tree T_Γ , whose projection in Γ is of surface type, is also said to be of surface type. The surfaces corresponding to surface type vertices of Γ are called the *surfaces of Γ* .

Remark 4.13: Note that the choice of the set V_S is part of the structure: it does not necessarily contain all the vertices of Λ which satisfy the conditions listed above.

Definition 4.14: (surface type automorphism) Let G be a group which admits a decomposition as a graph of groups with surfaces Λ , and let S be a surface type vertex group in this decomposition. An automorphism of S which restricts to conjugation by an element of S on each maximal boundary subgroups has a standard extension (recall section 3.1) to an automorphism of the whole group, which we call a *surface type automorphism*.

It is a famous result, proved first by Dehn and later independently by Lickorish (see [Lic64]), that if S is the fundamental group of an orientable surface with boundary, the group of automorphisms of S which preserve the conjugacy class of boundary subgroups is generated by Dehn twists of splittings of S in which boundary subgroups are elliptic. Thus in the orientable case, surface type automorphisms corresponding to S are in the subgroup generated by Dehn twists of G .

Definition 4.15: (modular group $\text{Mod}(\Lambda)$ of a graph of groups Λ) Let G be a group which admits a decomposition as a graph of group with surfaces Λ . The modular group $\text{Mod}(\Lambda)$ of Λ is the subgroup of $\text{Aut}(G)$ generated by inner automorphisms, Dehn twists, generalised Dehn twists, and surface type automorphisms.

Definition 4.16: (abelian modular group $\text{Mod}(G)$ of a group G) Let G be a finitely generated group. We define the abelian modular group of G , denoted by $\text{Mod}(G)$, to be the subgroup of $\text{Aut}(G)$ generated by the modular groups of all the abelian splittings of G .

4.2.2 Action shortening result

Recall that in definition 3.25, we defined the length of an action λ of G on a pointed metric space (X, x) to be the maximal displacement of the basepoint x by an element of the generating set Σ_G . If σ is an automorphism of G , we denote $\lambda \circ \sigma$ the action of G on (X, x) given by $(g, x) \mapsto \lambda(\sigma(g), x)$, and we give

Definition 4.17: (short action) An action λ of a group G on a pointed space (X, x) is short if for any element σ of $\text{Mod}(G)$, the length of λ is at most the length of $\lambda \circ \sigma$.

The action shortening result we want to state now is a slightly altered version of the one proved by Rips and Sela (see [RS94] or [Sel97a]). It asserts that, under the right set of conditions, if a sequence λ_n of actions of a finitely generated freely indecomposable group G on pointed hyperbolic spaces (X_n, x_n) converges (in the equivariant Gromov Hausdorff topology) to an action λ on a real tree T , then at most finitely many of the actions λ_n are short.

There are various possible sets of hypotheses on the G -spaces X_n and on the limit G -tree T for which some shortening result holds. The hypotheses on T should enable us to analyse it using Rips theory, which decomposes real G -trees into simple building blocks of given types (see Theorem 10.8 in [RS94], or Theorem 5.1 of [Gui08]). One of the conditions an action needs to satisfy for Rips theory to apply is the following.

Definition 4.18: (superstable) An action on a real tree is said to be *superstable* if for any pair of arcs J, I with $J \subseteq I$ and $\text{Fix}(I) \neq 1$, we have $\text{Fix}(I) = \text{Fix}(J)$.

In Theorem 10.8 of [RS94], Rips and Sela give the existence of a decomposition for a real tree under a weaker condition, however Guirardel showed in [Gui08] that this stronger hypothesis is necessary. To see that an action is superstable, we will use the following criterion.

Lemma 4.19: *If a group G acts on a real tree T in such a way that any subgroup fixing a tripod is trivial, and any subgroup fixing an arc is abelian, then the action is superstable.*

Proof. Indeed, let $I = [a, b]$ and $J = [c, d]$ be two non-trivial arcs of T with $J \subseteq I$. We clearly have $\text{Fix}(I) \leq \text{Fix}(J)$. If we do not have equality, there is an element g which lies in $\text{Fix}(J)$ but not in $\text{Fix}(I)$ so without loss of generality $g \cdot b \neq b$. We now want to see that $\text{Fix}(I)$ must be trivial. Let $h \in \text{Fix}(I)$, and note that g and h commute since they both fix J . We have $h \cdot a = a$ and $h \cdot b = b$ since $h \in \text{Stab}(I)$, and $h \cdot (g \cdot b) = gh \cdot b = g \cdot b$. The element h fixes the tripod formed by a , b and $g \cdot b$, thus it must be trivial. \square

Once we know that the limit tree can be decomposed by Rips' analysis, we need to add some conditions to deal with the different types of building blocks. We require for example that G be freely indecomposable to ensure that there are no Levitt (or thin) components. To be able to deal with axial components, the assumption that solvable subgroups are free abelian will prove useful.

As for the part of the shortening argument which deals with simplicial components of the limit tree, they require some hypotheses on the G -spaces X_n . The hypotheses must be sufficient to give some understanding of how an element which fixes an arc in the limit action acts on an approximation of this arc in X_n for n large enough.

If the spaces X_n are trees, this is much easier to achieve, since isometries of trees are very easily described. In this case, if an element fixes an arc I in the limit action, it is easy to see that it must act in X_n as a translation whose axis contains a segment approximating I . It will be important to know that this translation is not trivial: a strong hypothesis which ensures this is the assumption that the actions are all free, or that for each g in G , the action of g on X_n is free for n large enough. This will be satisfied if the spaces X_n are rescaling of spaces of the form $X[h_n]$ for some sequence of morphisms h_n to a free group whose stable kernel is trivial. Such a set of conditions is used to show that limit groups admit factors sets (see [Sel01], or [Wil06]). Another possibility is to assume that the diameter of the fixed point set of an element is bounded by d_n , with d_n tending to 0. This is the case if the X_n are rescaled k -acylindrical G -trees for example, as is used in [Sel97a].

If the spaces X_n are not trees, but only δ_n -hyperbolic spaces, we have to be slightly more careful. In [RS94], for example, the authors assume that all the actions are free, and that the number of translates of a point p which are at a distance at most $10\delta_n$ of p is bounded uniformly in n and in p .

However, if we know that the actions λ_n are rescalings of actions of the form $X[h_n]$, where h_n is a morphism into a δ -hyperbolic group, the proof is greatly simplified. Indeed in this case, X_n is proper and geodesic, and for any element g of G which is not in the stable kernel, for n large enough, the action of g on X_n is hyperbolic. Moreover given a non-trivial element g of G , there is a fixed power of g which has translation length greater than $12\delta_n$ in all the spaces X_n for n large enough (recall lemma 3.14), and this will also prove useful.

We can now state the action shortening theorem

Theorem 4.20: *Let G be a torsion-free and freely indecomposable group, endowed with a finite generating set Σ_G . Suppose moreover that solvable subgroups of G are free abelian groups.*

Let $(X_n, x_n)_{n \in \mathbb{N}}$ be a sequence of pointed proper and geodesic δ_n -hyperbolic metric space endowed with actions λ_n of G by isometries. Suppose that the sequence $(X_n, x_n)_{n \in \mathbb{N}}$ converges to a pointed real G -tree (T, x) . Assume that any non-trivial arc stabiliser of T contains an element which, for all n large enough, acts hyperbolically on X_n , with translation length at least $12\delta_n$. If the action λ of G on T satisfies:

1. λ is non-trivial;
2. tripod stabilisers are trivial;

3. pointwise arc stabilisers are abelian;

4. λ is superstable;

then for n large enough, the actions λ_n are not short.

4.2.3 Morphism shortening result

Let Γ be a torsion-free hyperbolic group endowed with a finite generating set Σ_Γ . Denote by X its Cayley graph. In lemma 4.2, we saw that if we choose basepoints properly, we can build from a sequence of morphisms $h_n : G \rightarrow \Gamma$ a sequence of G -spaces converging to a real G -tree. We now want to find sufficient conditions on the morphisms h_n to ensure that the sequence and its limit satisfy the hypotheses of the action shortening result we just saw.

We first need to choose basepoints for the spaces $X[h_n]$ so that the limit of our sequence of G -spaces is a real tree endowed with a non-trivial action of G .

Definition 4.21: (minimal displacement, minimally displaced point) *Let G, Γ be groups endowed with finite generating sets Σ_G and Σ_Γ respectively. Let $h : G \rightarrow \Gamma$ be a morphism. The minimal displacement $\mu[h]$ is the infimum of the function $\Gamma \rightarrow \mathbb{N}$ given by*

$$x \mapsto \max_{s \in \Sigma_G} |x^{-1}h(s)x|_{\Sigma_\Gamma}.$$

Since the word metric is integer valued, this infimum is a minimum. The point where it is reached is called the minimally displaced point of h , and denoted by $x[h]$. We will slightly abuse notations and identify $x[h]$ to the corresponding point of the Cayley graph X of Γ with respect to Σ_Γ .

We have $\mu[h] = \max_{s \in \Sigma_G} d_X(x[h], h(s) \cdot x[h])$, so the length of the action of G on $(X[h], x[h])$ is precisely the minimal displacement $\mu[h]$ of h .

Definition 4.22: (short morphism) *A morphism $G \rightarrow \Gamma$ is short if for any element σ of $\text{Mod}(G)$, we have*

$$\mu[h] \leq \mu[h \circ \sigma]$$

In other words, the morphism h is short if and only if the action of G on $(X[h], x[h])$ is short. We can then show

Proposition 4.23: *Let G be a torsion-free group, endowed with a finite generating set Σ_G , all of whose virtually abelian subgroups are abelian. Let Γ be a torsion-free hyperbolic group endowed with a finite generating set Σ_Γ .*

Let $h_n : G \rightarrow \Gamma$ be a sequence of pairwise non-conjugate morphisms, and suppose there is a non-principal ultrafilter ω for which $\text{Ker}_\omega(h_n)$ is trivial.

Then the ultraproduct of the spaces $\frac{1}{\mu[h_n]}(X[h_n], x[h_n])$ with respect to ω is defined, and it is a real tree (T, x) which satisfies the conditions 1 to 4 in 4.20.

Proof. To see that the ultraproduct of the spaces $\frac{1}{\mu[h_n]}(X[h_n], x[h_n])$ with respect to ω is well defined and is a real tree, it is enough to show by remark 4.2 that $\mu[h_n]$ tends to infinity with respect to ω . Let h be a morphism $G \rightarrow \Gamma$. Note that for every g in Σ_G , we have

$$|x^{-1}[h]h(g)x[h]|_{\Sigma_G} = d_X(1, x^{-1}[h]h(g)x[h]) = d_X(x[h], h(g)x[h]) \leq \mu[h].$$

Thus h has a conjugate which sends all the generators of G in the ball of radius $\mu[h]$ in Γ . There are only finitely many such morphisms, so since the h_n are pairwise non-conjugate, $\mu[h_n]$ must tend to infinity with respect to ω .

Suppose that the action is trivial. If $y = (y_n)$ is a point in $T = X[h, \omega]$, for every g in Σ_G we have $\lim_\omega d_n(y_n, g \cdot y_n) = 0$. Thus there exists A_g in ω such that, for any n in A_g , we have

$d_n(y_n, g \cdot y_n) < 1/2$. For n in $\bigcap_{g \in \Sigma_G} A_g$, we have $\max_{g \in \Sigma_G} d_n(y_n, g \cdot y_n) < 1/2$. By definition of $\mu[h_n]$, we have

$$\begin{aligned} \max_{g \in \Sigma_G} d_n(y_n, g \cdot y_n) &= \frac{1}{\mu[h_n]} \max_{g \in \Sigma_G} d_X(y_n, h_n(g) \cdot y_n) \\ &\geq \frac{1}{\mu[h_n]} \max_{g \in \Sigma_G} d_X(x[h_n], h_n(g) \cdot x[h_n]) \\ &\geq 1 \end{aligned}$$

This gives a contradiction. Thus the action of G on T is non-trivial.

By lemma 4.5, we see that elements fixing a non-trivial tripod must be trivial. Since G is torsion-free, in fact the tripods stabilisers themselves are trivial. By lemma 4.7, the pointwise stabiliser of a non-trivial arc is abelian. By remark 4.19, the action is also superstable. \square

We can now show the morphisms shortening results. We define

Definition 4.24: (stable sequence, stable kernel) *A sequence of morphisms $h_n : G \rightarrow \Gamma$ is said to be stable if for any element g , either $h_n(g)$ is trivial for all but finitely many n , or $h_n(g)$ is non-trivial for all but finitely many n . The stable kernel $\underline{\text{Ker}}(h_n)$ of such a sequence is the set of elements g of G for which the first alternative holds.*

Equivalently, a sequence is stable if and only if its stable kernels with respect to any two non-principal ultrafilters are equal.

Note that if G is finitely generated, any sequence of morphism $h_n : G \rightarrow \Gamma$ contains a stable subsequence. To see this, note that for any finite subset B of G , we can extract a subsequence of morphisms whose kernels all have the same intersection with B . Now let B_k be an exhausting sequence of finite subsets of G , and for each k , extract by induction a subsequence $(h_n^k)_{n \in \mathbb{N}}$ of $(h_n^{k-1})_{n \in \mathbb{N}}$. The diagonal subsequence $(h_n^n)_{n \in \mathbb{N}}$ is then stable.

Theorem 4.25: *Let G be a torsion-free and freely indecomposable group, endowed with a finite generating set Σ_G . Suppose that virtually solvable subgroups of G are free abelian. Let Γ be a torsion-free hyperbolic group endowed with a finite generating set Σ_Γ .*

Let $h_n : G \rightarrow \Gamma$ be a stable sequence of non-injective short morphisms. Then the stable kernel of $(h_n)_{n \in \mathbb{N}}$ is non-trivial.

Proof. Suppose by contradiction that $\underline{\text{Ker}}(h_n)$ is trivial. If the maps h_n belonged to finitely many conjugacy classes, there would be only finitely many possibilities for the kernel of h_n . Then h_n would admit a subsequence all of whose terms have the same kernel. Since $\underline{\text{Ker}}(h_n) = \{1\}$, this kernel would have to be trivial, which contradicts the non-injectivity of the maps h_n . Thus, up to extracting a subsequence, the maps h_n are pairwise non-conjugate.

Let ω be a non-principal ultrafilter. The stable kernel of $(h_n)_{n \in \mathbb{N}}$ with respect to ω is trivial, the morphisms are pairwise non-conjugate, and virtually abelian subgroups of G are solvable so they are abelian by hypothesis. We can thus apply proposition 4.23 to see that the ultraproduct with respect to ω of the spaces $\frac{1}{\mu[h_n]}(X[h_n], x[h_n])$ is a pointed G -tree T , which satisfies the conditions 1 to 4 in 4.20. Thus some subsequence of $(\frac{1}{\mu[h_n]}(X[h_n], x[h_n]))_{n \in \mathbb{N}}$ converges to a pointed G -tree T which satisfies the conditions 1 to 4 in 4.20.

Let g be a non-trivial element of G which fixes an arc in T : for all n large enough, $h_n(g)$ is non-trivial, so by lemma 3.14, for all n large enough, $h_n(g^{M_{12\delta}})$ has translation length at least 12δ in $X[h_n]$, so it has translation length at least $12\delta_n$ in $\frac{1}{\mu[h_n]}(X[h_n], x[h_n])$. The element $g^{M_{12\delta}}$ fixes the same arc as g , so the condition in 4.20 about stabilisers of arcs holds.

We can thus apply Theorem 4.20 to see that for all n large enough, the action of G on $\frac{1}{\mu[h_n]}(X[h_n], x[h_n])$ is not short. This contradicts the shortness of the maps h_n . \square

If we restrict ourselves to injective maps we get

Theorem 4.26: *Let G be a torsion-free and freely indecomposable group, endowed with a finite generating set Σ_G . Suppose that virtually solvable subgroups of G are free abelian. Let Γ be a torsion-free hyperbolic group endowed with a finite generating set Σ_Γ .*

There exists a finite set $\{i_1, \dots, i_k\}$ of embeddings $G \hookrightarrow \Gamma$ such that for any embedding $i : G \hookrightarrow \Gamma$, there is an index j with $1 \leq j \leq k$, an element γ of Γ , and a modular automorphism σ of G such that

$$i = \text{Conj}(\gamma) \circ i_j \circ \sigma.$$

Proof. If this were not the case, there would be an infinite sequence of pairwise non-conjugate short embeddings $h_n : G \rightarrow \Gamma$. The stable kernel of such a sequence with respect to any ultrafilter ω is of course trivial. Up to extraction, and by proposition 4.23, the sequence $(\frac{1}{\mu[h_n]}(X[h_n], x[h_n]))_{n \in \mathbb{N}}$ converges to a pointed G -tree T which satisfies the conditions 1 to 4 in Theorem 4.20. Thus for n large enough, these actions are not short: this contradicts the shortness of the embeddings h_n . \square

4.3 Shortening morphisms in the relative case

We will also use a relative version of the shortening argument. Here relative means that we fix a subgroup H of a group G , and we ask that this subgroup fixes a point in the actions of G on trees (real or simplicial). Similarly, instead of asking G to be freely indecomposable, we require that it be freely indecomposable relative to H (i.e. that no proper free factor of G contains H). Apart from a few modifications of the sort, the arguments are very similar to the non-relative case. The main difference that should be noted lies in the proof of non-triviality of the limit action in the proof of proposition 4.31, compared to that found in the proof of proposition 4.23.

4.3.1 Action shortening result

We start by adapting our definition of the modular group.

Definition 4.27: (relative abelian modular group $\text{Mod}_H(G)$) *Let G be a finitely generated group, and let H be a subgroup of G . Let Λ be a splitting of G as a graph of groups with surfaces for which H is elliptic. The modular group $\text{Mod}_H(\Lambda)$ of Λ relative to H is the subgroup of $\text{Mod}(\Lambda)$ containing all the automorphisms which fix H . The abelian modular group of G relative to H , denoted $\text{Mod}_H(G)$, is the subgroup of $\text{Aut}(G)$ generated by the subgroups $\text{Mod}_H(\Lambda)$, where Λ is an abelian splitting of G in which H is elliptic.*

Theorem 4.20 generalises to

Theorem 4.28: *Let G be a torsion-free group endowed with a finite generating set Σ_G , whose solvable subgroups are free abelian. Let H be a subgroup of G , and assume G is freely indecomposable relative to H .*

Let $(X_n, x_n)_{n \in \mathbb{N}}$ be a sequence of proper and geodesic pointed δ_n -hyperbolic spaces, endowed with isometric actions λ_n of G by isometries.

Suppose that the sequence $(X_n, x_n)_{n \in \mathbb{N}}$ converges to a pointed real G -tree (T, x) . Assume that any non-trivial arc stabiliser of T contains an element which, for all n large enough, acts hyperbolically on X_n with translation length at least $12\delta_n$. If the action λ of G on T satisfies:

1. λ is non-trivial;
2. tripod stabilisers are trivial;
3. pointwise arc stabilisers are abelian;

4. λ is superstable;

5. H fixes a point;

then for n large enough, the actions λ_n are not short.

4.3.2 Morphism shortening result

Let G be a group endowed with a finite generating set Σ_G . Let Γ be a torsion-free hyperbolic group endowed with a finite generating set Σ_Γ . Denote by X the Cayley graph of Γ with respect to Σ_Γ . Let H be a subgroup of G , and fix an embedding of H in Γ so that H is also a subgroup of Γ .

Here again, we want to find sufficient conditions on a sequence of morphisms $h_n : G \rightarrow \Gamma$, so that we can build from it a sequence of actions which satisfies the hypotheses of the action shortening result we just saw.

The main difference is that we need H to fix a point in the limit. This will affect our choice of basepoint: we will not choose as a basepoint the minimally displaced point, but simply the point e of X corresponding to the identity in Γ . To make sure H fixes the basepoint in the limit, we will require that the following condition hold.

Definition 4.29: (fixing H in the limit) *We say that a sequence of morphisms $h_n : G \rightarrow \Gamma$ fixes H in the limit if for all $r \in \mathbb{N}$, for n large enough, h_n is the identity on words of length less or equal to r (i.e. on the finite set $B_G(r) \cap H$).*

Remark 4.30: *If H is non-abelian, and if $(h_n)_{n \in \mathbb{N}}$ is a sequence of pairwise distinct maps which fixes H in the limit, the maps h_n are pairwise non-conjugate for n large enough. Indeed, let $a_1, a_2 \in H$ such that $[a_1, a_2] \neq 1$. For all n greater than some constant n_0 , $h_n(a_1) = a_1$ and $h_n(a_2) = a_2$ so that if γ is a non-trivial element of Γ , it cannot commute both with $h_n(a_1)$ and $h_n(a_2)$ (recall that torsion-free hyperbolic groups are commutative-transitive). Thus without loss of generality, $\text{Conj}(\gamma) \circ h_n(a_1) \neq a_1$, and $\text{Conj}(\gamma) \circ h_n \neq h_r$ for $r \geq n_0$.*

The following proposition expresses sufficient conditions on a sequence of morphisms $h_n : G \rightarrow \Gamma$ to enable us to apply Theorem 4.28.

Proposition 4.31: *Let G be a group endowed with a finite generating set Σ_G whose virtually abelian subgroups are abelian. Let Γ be a torsion-free hyperbolic group endowed with a finite generating set Σ_Γ . Let H be a non-abelian subgroup of G with a fixed embedding in Γ .*

Let $h_n : G \rightarrow \Gamma$ be a sequence of pairwise distinct morphisms which fixes H in the limit, and suppose that ω is a non-principal ultrafilter such that $\text{Ker}_\omega(h_n)$ is trivial.

Then the ultraproduct of the spaces $\frac{1}{l(X[h_n], e)}(X[h_n], e)$ with respect to ω is a real G -tree which satisfies the conditions 1 to 5 in 4.28.

Proof. The fact that the h_n are pairwise distinct implies that the maps h_n send the elements of Σ_G outside balls of larger and larger radius. Recall that

$$l(X[h_n], e) = \max_{g \in \Sigma_G} d_X(e, h_n(g) \cdot e) = \max_{g \in \Sigma_G} |h_n(g)|_{\Sigma_\Gamma},$$

so the sequence $l(X[h_n], e)$ tends to infinity. By lemma 4.2, the ultraproduct of the spaces $\frac{1}{l(X[h_n], e)}(X[h_n], e)$ with respect to ω is a pointed real G -tree T . We denote the basepoint of T also by e .

By lemma 4.5, subgroups of G fixing tripods are trivial, and as G is torsion-free, tripod stabilisers are trivial. Lemma 4.7 implies that non-trivial arcs have trivial pointwise stabilisers. By remark 4.19, the action is superstable.

For an element a of H , the image $h_n(a)$ is constant and equal to a for n large enough, so

$$\begin{aligned} d_\omega(e, a \cdot e) &= \lim_{\omega} d_n(e, a \cdot e) \\ &= \lim_{\omega} (d_X(e, h_n(a) \cdot e) / l(X_n, e)) \\ &= \lim_{\omega} |a|_{\Sigma_\Gamma} / l(X_n, e) = 0. \end{aligned}$$

Thus H fixes the point e .

Suppose that e is a global fixed point: then for each g in Σ_G , the set A_g of indices n for which $d_n(e, g \cdot e) < 1/2$ lies in ω . Thus for $n \in \bigcap_{g \in \Sigma_G} A_g$, we have $\max_{g \in \Sigma_G} d_n(e, g \cdot e) < 1/2$. By definition of $l(X[h_n], e)$ we have

$$\max_{g \in \Sigma_G} d_n(e, g \cdot e) = \frac{\max_{g \in \Sigma_G} d_X(e, h_n(g) \cdot e)}{l(X[h_n], e)} = 1$$

This is a contradiction, so e is not a global fixed point.

But now, if T did admit a global fixed point y , the non-trivial path between y and e would be fixed by H . Since pointwise arc stabilisers are abelian and H isn't, the action of G on T is non-trivial. \square

The change in our choice of basepoint means that our definition of a short morphism must also be slightly modified.

Definition 4.32: (short morphism with respect to H) *A morphism $h : G \rightarrow \Gamma$ is short with respect to H if for any element σ of $\text{Mod}_H(G)$, we have*

$$\max_{g \in \Sigma_G} d_X(e, h(g) \cdot e) \leq \max_{g \in \Sigma_G} d_X(e, h \circ \sigma(g) \cdot e).$$

In other words, h is short with respect to H if the action of G on $(X[h], e)$ is short. We can now show the relative morphism shortening result.

Theorem 4.33: *Let G be a torsion-free group, endowed with a finite generating set Σ_G . Suppose that the virtually solvable subgroups of G are free abelian. Let Γ be a torsion-free hyperbolic group endowed with a finite generating set Σ_Γ . Let H be a non-abelian subgroup of G , with respect to which G is freely indecomposable. Fix an embedding of H in Γ so that H is also a subgroup of Γ .*

Let $h_n : G \rightarrow \Gamma$ be a stable sequence of non-injective morphisms which fix H in the limit and are short with respect to H . Then the stable kernel of $(h_n)_{n \in \mathbb{N}}$ with respect to ω is non-trivial.

Proof. Suppose by contradiction that $\text{Ker}(h_n)$ is trivial. If $(h_n)_{n \in \mathbb{N}}$ has a constant subsequence, the maps in this subsequence must have trivial kernel, which contradicts their non-injectivity. Thus, $(h_n)_{n \in \mathbb{N}}$ has no constant subsequence, so up to extracting, we may assume that the maps h_n are pairwise distinct.

By proposition 4.31, up to another extraction, $(\frac{1}{l(X[h_n], e)}(X[h_n], e))_{n \in \mathbb{N}}$ tends to a real G -tree T which satisfies conditions 1 to 5 of Theorem 4.28.

Moreover, if g is a non-trivial element fixing an arc of T , we know by lemma 3.14 that for all n large enough, $h_n(g^{M_{12\delta}})$ has translation length at least 12δ .

Thus we can apply Theorem 4.28, which tells us that for n large enough, the action of G on $(X[h_n], e)$ is not short. This contradicts the shortness of the map h_n relative to H . \square

Similarly to the non-relative case, we can also give a version for injective maps.

Theorem 4.34: *Let G be a torsion-free group, endowed with a finite generating set Σ_G . Suppose that virtually solvable subgroups of G are free abelian groups. Let H be a non-abelian subgroup of G , with respect to which G is freely indecomposable. Let Γ be a torsion-free hyperbolic group*

endowed with a finite generating set Σ_Γ . Fix an embedding of H in Γ so that H is a subgroup of Γ .

There exists a finite set i_1, \dots, i_k of embeddings $G \hookrightarrow \Gamma$ such that for any embedding $i : G \hookrightarrow \Gamma$ which fixes H , there is an index j with $1 \leq j \leq k$, and an element σ of $\text{Mod}_H(G)$ such that

$$i = i_j \circ \sigma.$$

Remark 4.35: It looks like we get a much better result that in the non-relative case, since we got rid of the conjugation. However recall that since H is non-abelian, if a map h fixes H , a conjugate $\text{Conj}(\gamma) \circ h$ of h fixes H if and only if $\gamma = 1$.

Proof. Suppose $h_n : G \rightarrow \Gamma$ is a sequence of pairwise distinct embeddings fixing H . It is stable with a trivial stable kernel, so we can see by proposition 4.31 that a subsequence of the spaces $(\frac{1}{l(X[h_n], e)}(X[h_n], e))_{n \in \mathbb{N}}$ tends to a real G -tree T which satisfies conditions 1 to 5 of Theorem 4.28. Thus the action of G on $(X[h_n], e)$ is not short for n large enough, so that h_n is not short relative to H .

This shows that there is only a finite number of distinct short embeddings $G \rightarrow \Gamma$, which proves the result. \square

4.4 The relative Co-Hopf property for hyperbolic groups

We will prove a result which is a direct consequence of the shortening argument. It expresses the fact that torsion-free hyperbolic groups satisfy a relative co-Hopf property. Sela showed in [Sel97b] that freely indecomposable hyperbolic groups are co-Hopfian, but the proof is actually much harder than in the relative case.

Proposition 4.36: *Let G be a torsion-free hyperbolic group. Let H be a non-abelian subgroup of G relative to which G is freely indecomposable. If $\phi : G \rightarrow G$ is injective and fixes H then it is an isomorphism.*

Proof. Suppose ϕ is a strict embedding: then $\phi^n(G)$ is a strictly decreasing sequence of subgroups of G which are all isomorphic to G by isomorphisms fixing H .

The group G is torsion-free hyperbolic, so it satisfies all the hypotheses of proposition 4.34. As a consequence, the number of subgroups of G isomorphic to G by isomorphisms fixing H is finite. This is a contradiction. \square

Now we can actually get a stronger statement by using the following lemma, whose proof was suggested by Vincent Guirardel.

Lemma 4.37: *If a finitely generated group G is freely indecomposable relative to a subgroup H , then H has a finitely generated subgroup H_0 relative to which G is freely indecomposable.*

Proof. Suppose G' is a subgroup of G . Denote by $T(G')$ the set of all simplicial G -trees τ with trivial edge stabilisers in which G' fixes a vertex v_τ . Define

$$A(G') = \bigcap_{\tau \in T(G')} \text{Stab}(v_\tau)$$

To each τ in $T(G')$, we associate the corresponding free product decomposition of G . The number of factors of such a decomposition is bounded, since G is finitely generated: let $m_G(G')$ be the maximal number of factors that such a decomposition can have. A decomposition with $m_G(G')$ factors is clearly of the form

$$A * B_1 * \dots * B_r$$

where B_1, \dots, B_r are freely indecomposable (possibly cyclic), and A contains G' and is freely indecomposable with respect to G' . Now $A(G') < A$ since this decomposition corresponds to a tree τ in $T(G')$ for which $\text{Stab}(v_\tau) = A$. But A is freely indecomposable with respect to G' , thus in any tree τ of $T(G')$, A fixes the vertex v_τ . Thus $A = A(G')$. Note that this implies that $A(G')$ is a free factor of G , and is freely indecomposable with respect to G' .

If $G'' < G'$, then $T(G') \subseteq T(G'')$. This implies that $A(G'') < A(G')$, and $m_G(G') \leq m_G(G'')$. Moreover, if $m_G(G') = m_G(G'')$, a maximal decomposition for G' is also maximal for G'' , thus $A(G'') = A(G')$.

We can now prove the lemma. Let $\{h_1, h_2, \dots\}$ be a generating set for H . Consider the sequence of subgroups $H_k = \langle h_1, \dots, h_k \rangle$ of H . By the remarks above, the sequence $(m_G(H_k))_{k>0}$ is decreasing, and bounded below by 1. Thus it must stabilise, and by what we saw above, this implies that the sequence $A(H_k)$ stabilises after some index k_0 . In particular $H_k < A(H_{k_0})$ for all k , so $H < A(H_{k_0})$. But $A(H_{k_0})$ is a free factor of G : since we assumed G freely indecomposable with respect to H , we must have $A(H_{k_0}) = G$, so G is freely indecomposable with respect to H_{k_0} . \square

Proposition 4.38: *Let G be a torsion-free hyperbolic group. Let H be a non-cyclic subgroup of G . Suppose G freely indecomposable relative to H . There exists a finite subset F_0 of H such that if $\phi : G \rightarrow G$ is an injective morphism which fixes F_0 , then it is an isomorphism.*

Proof. Just take F_0 to be a generating set for the subgroup H_0 given by lemma 4.37. If ϕ fixes F_0 , it fixes H_0 relative to which G is freely indecomposable. Thus we can apply proposition 4.36 to G with the subgroup H_0 , to deduce that ϕ is an isomorphism. \square

Chapter 5

Proof of the action shortening Theorem

We will present a proof of theorem 4.28. The strategy is as follows: we start by analysing the G -tree T using Rips theory. Rips theory enables us, under certain hypotheses like superstability, to analyse actions of finitely generated groups on real trees by decomposing such an action into elementary building blocks (see [Sel97a] and [Gui08]). Then we produce for each type of blocks an automorphism of G which will shorten all the paths of the form $[x, u \cdot x]$ which intersect one of these blocks, where u is an element of the generating set.

The proof we give is based on the proof of Theorem 4.3 of [RS94], the fact that we are in the relative case does not introduce particular difficulties. However we altered the presentation, mainly when dealing with the discrete case, and in general we give a more detailed version of the various arguments. We followed also the proof given in [Wil06], where the non-relative version of the theorem is proved in the specific case where the actions come from a sequence of homomorphisms into the free group.

5.1 Some examples of actions on real trees

Let us first give some classical examples of actions on real trees.

Example 5.1: (simplicial type) Let T be a real G -tree. Branching points in a real tree are points whose complement has more than two connected components. If branching points are isolated, we say that T is a simplicial G -tree, or that the action is of simplicial type.

Example 5.2: (axial type) Let T be a line, and let G act on T with dense orbits, in such a way that the image of G in $\text{Isom}(\mathbb{R})$ is finitely generated. We say that the action of G on T is of axial type. We then have the following exact sequence for G :

$$1 \rightarrow K \rightarrow G \rightarrow H \rightarrow 1,$$

where K fixes T pointwise, and H is a finitely generated subgroup of $\text{Isom}(\mathbb{R})$ of rank at least 2.

Example 5.3: (surface type) Consider a foliation \mathcal{F} endowed with a transverse measure μ on a surface with boundary Σ , as defined in [FLP79] (see also [LP97] or section 1.7 of [Gui08]). For $x \in \Sigma$, let O_x be the set of points y such that there is a path $[x, y]$ which is contained in a leaf, and which meets no singularity of the foliation. Suppose that for any point x of Σ , the set O_x is dense in Σ (a foliation which satisfies this is called arational). The foliation \mathcal{F} lifts to a foliation

$\tilde{\mathcal{F}}$ of a universal cover $\tilde{\Sigma}$ of Σ , and μ induces a pseudometric on the set of leaves of $\tilde{\mathcal{F}}$ given by

$$d(x, y) = \inf_{\gamma} \mu(\gamma)$$

where γ ranges over all the paths in Σ which lift to a path from x to y . Quotienting the space of leaves by the equivalence relation induced by this pseudometric, we get a real tree T (see [MS91], where laminations are used instead of foliations, or [LP97]). We then say that the action of the fundamental group S of Σ on T is of surface type. It is a result of [MS91] that such an action is free if Σ has no boundary, and it can be generalised to show that in the case with boundary, the action on the dual tree has trivial arc stabilisers, and that the only elliptic elements are the boundary elements of S .

A way to get surfaces endowed with foliations is to give them a band complex structure over an interval. To do so, we need the following definitions (see [BF95], or definition 2.2 of [Wil]).

Definition 5.4: (union of bands) *Let I be an interval of the real line. Given a finite set of pairs (I_i, I'_i) of subintervals of I , and isometries $\psi_i : I_i \rightarrow I'_i$, we can build a topological space Y by gluing for each i the band $I_i \times [0, 1]$ to I via the isometries*

$$\left\{ \begin{array}{ll} I_i \times \{0\} & \rightarrow I_i \\ (x, 0) & \mapsto x \end{array} \right. \quad \text{and} \quad \left\{ \begin{array}{ll} I_i \times \{1\} & \rightarrow I'_i \\ (x, 1) & \mapsto \psi_i(x) \end{array} \right.$$

Such a space Y is called a union of bands on I .

Definition 5.5: (union of bands of surface type) *A union of bands Y on an interval I for which all but finitely many points of I lie in exactly two bands is said to be of surface type.*

Suppose Y is a union of bands of surface type. Topologically, Y is a surface with boundaries, and its fundamental group is the free group generated by elements $(g_i)_{1 \leq i < n}$ corresponding to the loops given by the various bands.

A band $I_i \times [0, 1]$ is foliated by sets of the form $\{x\} \times [0, 1]$. This foliation admits a standard transverse measure, where the measure of a transverse arc γ is the length of its projection on the interval I_i . The union of bands Y thus admits a natural measured foliation induced by this foliation of the bands. Note that every boundary component of Y lies in some leaf of Y .

Definition 5.6: (complex of bands) *The space X is said to be a complex of bands of surface type with underlying union of bands Y if*

- Y is a union of bands of surface type,
- X is obtained from Y by gluing 2-cells along some of the boundary components of Y .

Note that X is also a surface with (possibly empty) boundary, and $\pi_1(X)$ is generated by the elements g_i corresponding to the bands of Y . The relations satisfied by these elements are words $w(g_1, \dots, g_n)$ formed by following the boundary of a 2-cell and reading out the name of the bands whose boundary we follow. We call this presentation of $\pi_1(X)$ the presentation associated to X .

Thus X is a surface endowed with a measured foliation. We denote by T_X the $\pi_1(X)$ -tree associated to the foliation on X as described in example 5.3. We can choose a lift of I in the universal cover $(\tilde{X}, \tilde{\mathcal{F}})$, it is transverse to $\tilde{\mathcal{F}}$ so this gives us an injection of I in T_X .

In fact, it can be shown that any arational foliation \mathcal{F} on a surface with boundary Σ endowed with a transverse measure μ can be built in this way: Σ admits a band complex structure X whose natural measured foliation is homotopy equivalent to (\mathcal{F}, μ) . Equivalently, for any surface type action of S on a tree T , there is a band complex X such that T and T_X are isomorphic as real S -trees.

The idea of the proof is to pick on the surface an interval transverse to the foliation, and cut along leaves which contain singular points of the foliation until we meet the chosen interval. If the singular point was not contained in a boundary, we glue a 2-cell to the boundary thus created.

Note that by the description above, a structure of band complex X on Σ gives us a presentation for S .

Remark 5.7: Recall that the base interval I of the union of bands Y embeds in T_X . The generators of the presentation associated to the band complex structure X on P translate any point x of I by a distance smaller than the length of the interval I .

5.2 Graphs of actions

The notion of graph of actions allows us to describe the decomposition of a real G -tree into various components. They were introduced by Levitt in [Lev94], but we present the slightly different definition given in [Gui08]:

Definition 5.8: (graph of actions) Let G be a group. A graph of actions on R -trees is given by $\mathcal{G} = (\tau, (T_v)_{v \in V(\tau)}, (p_e)_{e \in E(\tau)})$ where

1. τ is a simplicial G -tree called the skeleton;
2. for each vertex v of τ , T_v is a real tree called the vertex tree;
3. for each edge e of τ , p_e is a point of $T_{t(e)}$ called the attaching point of e .

Moreover we require that the following equivariant properties be respected

1. G acts on the disjoint union of the trees T_v in such a way that the map $T_v \rightarrow v$ is equivariant;
2. for $g \in G$ and $e \in E(\tau)$, we have $p_{g \cdot e} = g \cdot p_e$.

We associate to \mathcal{G} the G -tree $T_{\mathcal{G}}$ obtained by quotienting the disjoint union of the T_v by the relations $p_e \sim p_{\bar{e}}$. The trees T_v inject in $T_{\mathcal{G}}$, their images are called components of T , and still denoted by T_v .

We say that a G -tree T splits as a graph of actions if T is isomorphic to $T_{\mathcal{G}}$ for some graph of actions \mathcal{G} .

Remark 5.9: The definition implies that T_v is invariant under the action of the stabiliser G_v of v in τ , and that p_e is invariant under the action of the stabiliser G_e of e in τ .

The following result expresses, in a tree which splits as a graph of actions, the distance between a point and one of its translates. It will prove very useful in the sequel.

Let \mathcal{G} be a graph of actions, and let Δ_{τ} be the graph of groups corresponding to the action of G on τ . Denote by π the quotient map $\tau \rightarrow G \backslash \tau$. Choosing a fundamental domain τ_0 in τ and Bass-Serre elements t_e for each edge e in τ_0 gives us an isomorphism between G and the fundamental group of Δ_{τ} . According to this isomorphism, an element of G can be written as a word in the elements t_e and in the elements of the groups G_v . Moreover, given a vertex $\pi(v)$ of Δ_{τ} , we can always choose to represent g by a word of the form

$$g_1 t_{e_1} g_2 t_{e_2} \dots t_{e_l} g_{l+1},$$

where the path formed by the edges $\pi(e_1), \dots, \pi(e_l)$ forms a loop based at $\pi(v)$, the element g_i lies in $G_{o(\pi(e_i))}$, the element g_{l+1} lies in $G_{\pi(v)}$, and if $\pi(e_{i+1}) = \pi(\bar{e}_i)$, the element g_i is not trivial. We call this the loop representation of g based at $\pi(v)$ (it depends on the choice of a fundamental domain and Bass-Serre elements).

Lemma 5.10: *Let v be a vertex of our fundamental domain τ_0 , and let $x \in T_v$. Let g be an element of G , whose loop representation based at $\pi(v)$ with respect to our choice of fundamental domain and Bass-Serre elements is $g_1 t_{e_1} g_2 t_{e_2} \dots t_{e_l} g_{l+1}$. Then the path $[x, g(x)]$ is the concatenation of the following arcs*

- $I_0 = [x, g_1 \cdot p_{e_1}]$;
- $I_i = g_1 t_{e_1} \dots g_i t_{e_i} \cdot [t_{e_i}^{-1} \cdot p_{e_i}, g_{i+1} \cdot p_{e_{i+1}}]$ for $1 \leq i < l$;
- $I_l = g_1 t_{e_1} \dots g_l t_{e_l} \cdot [t_{e_l}^{-1} \cdot p_{e_l}, g_{l+1} \cdot x]$.

so that we have in T_g

$$d(x, g \cdot x) = d(x, g_1 \cdot p_{e_1}) + \sum_{i=1}^{l-1} d(t_{e_i}^{-1} \cdot p_{e_i}, g_{i+1} \cdot p_{e_{i+1}}) + d(t_{e_l}^{-1} \cdot p_{e_l}, g_{l+1} \cdot x).$$

Proof. The concatenation of the arcs I_0, \dots, I_l forms a path between x and $g \cdot x$. To see that it is indeed an arc, it is enough to show that there is no overlap. By the way we defined the loop representation, no two non-trivial arcs I_i, I_j lie in the same component of T : there can be no overlap. \square

5.3 Rips decomposition

We can now state the result of Rips theory we will need: it is essentially Theorem 3.1 of [Sel97a], except we have replaced the stability assumption by that of superstability, and we assume that G is torsion-free and freely indecomposable with respect to a subgroup H . We use the terminology developed in [Gui08], where Theorem 5.1 gives a generalisation of the result of Rips and Sela. Thus the following result can be seen as a particular case of Theorem 5.1 of [Gui08].

Theorem 5.11: *Let G be a finitely generated torsion-free group which is freely indecomposable with respect to one of its subgroups H . Suppose G acts minimally, non-trivially, and superstably on a real tree T by isometry. Suppose moreover that tripods are trivially stabilised, and that H fixes a point of T . Then T has a decomposition as a graph of actions $\mathcal{G} = (\tau, (T_v)_{v \in V(S)}, (p_e)_{e \in E(S)})$ where each vertex action is either*

1. *of simplicial type: T_v is a simplicial G_v -tree;*
2. *of surface type: T_v is dual to an arational measured foliation on a surface with boundary;*
3. *of axial type: T_v is a line, and the image of G_v in $\text{Isom}(T_v)$ is a finitely generated group which acts on it with dense orbits.*

Note that the assumption of trivial tripod stabilisers implies in particular that if T is not a line, the action is faithful.

Fix a generating set Σ_G for G . To prove theorem 4.28, we need to find for all n large enough an element σ_n of $\text{Mod}_H(G)$ such that the action $\lambda_n \circ \sigma_n$ is shorter than λ_n . For an element g of the generating set, consider the path $[x, g \cdot x]$ in the limit tree T .

Suppose that we managed to find an element σ of $\text{Mod}_H(G)$ such that for all g in Σ_G , the path $[x, \sigma(g) \cdot x]$ is strictly shorter than the path $[x, g \cdot x]$. Then this will suffice, since for n large enough, there will be an ϵ -approximation between T and X_n , with ϵ smaller than the difference of the lengths of $[x, g \cdot x]$ and $[x, \sigma(g) \cdot x]$, thus $d_n(x_n, \sigma(g) \cdot x_n) < d_n(x_n, g \cdot x_n)$.

We will see that we can find such a relative modular automorphism for paths $[x, g \cdot x]$ which intersect a surface or an axial type component (see Theorem 5.12 and Theorem 5.17). However, in the case where there are paths of the form $[x, g \cdot x]$ which lie completely in simplicial components, this will not be sufficient. For these, we will have to go to an approximation X_n of T (see Theorem 5.22), and the shortening modular automorphism we use will depend on n .

5.4 Surface case

The following theorem allows us to shorten paths that intersect a surface type component.

Theorem 5.12: *Let G be a finitely generated group acting on a real tree T which admits a decomposition as a graph of actions $\mathcal{G} = (\tau, (T_v)_{v \in V(\tau)}, (p_e)_{e \in E(\tau)})$. Denote by $\Delta_{\mathcal{G}}$ the graph of groups corresponding to the action of G on τ . Let U be a finite subset of G . There exists an element σ of $\text{Mod}(\Delta_{\mathcal{G}})$ such that for any element u of U ,*

- *if the geodesic segment $[x, u \cdot x]$ intersects some surface type components non-trivially, then*

$$d(x, \sigma(u) \cdot x) < d(x, u \cdot x);$$

- *if not, $\sigma(u) = u$.*

Let us first show the following lemma, which in particular implies Theorem 5.12 in the special case where the tree T consists of exactly one surface component.

Lemma 5.13: *Suppose that the fundamental group S of a surface with boundary Σ acts on a real tree T by an action of surface type. Then for any finite subset V of S , for any point z of T and for $\eta > 0$, there exists an automorphism ϕ of S which restricts to a conjugation on each boundary subgroup, and such that for any element v of V ,*

$$d(z, \phi(v) \cdot z) < \eta.$$

Before proving this, we define what we mean by a presentation of S as the fundamental group of a surface with boundary, and when we consider two such presentation to be equivalent.

Definition 5.14: (surface presentation) *Let S be the fundamental group of a surface with boundary Σ . A surface presentation P of S is given by a tuple (k, R, B, h) , where*

- *k is a positive integer;*
- *R and B are finite tuples of elements of $F_k = \langle a_1, \dots, a_k \rangle$, the free group of rank k ;*
- *h is a map $\{a_1, \dots, a_k\} \rightarrow S$ whose extension to F_k is surjective, has kernel normally generated by the elements of R , and sends the tuple B on a tuple of pairwise non-conjugate maximal boundary elements of S .*

To an element g of S we can associate the word length $|g|_P$ of g in the presentation P .

We say that two surface presentations (k, R, B, h) and (k', R', B', h') of S are combinatorially equivalent if $k = k'$, $R = R'$ and $B = B'$. Then, there is a natural automorphism of S given by sending $h(a_i)$ to $h'(a_i)$ for $1 \leq i \leq k$. It is clear that this isomorphism restricts to a conjugation on each boundary subgroup.

We can now prove lemma 5.13.

Proof. Let X be a structure of band complex for Σ over an interval I and with underlying union of bands Y , such that, as a real G tree, T_X is isomorphic to T . Let n be the number of bands in Y .

We want to show that we can modify X to get another band complex structure X' for (Σ, \mathcal{F}) , also of surface type, over a very small interval. This will give a presentation for S in which the generators have small translation length by remark 5.7.

Step 1: Consider the combinatorial equivalence classes of surface presentations of S for which $k \leq n+1$, $|R| \leq 2(n+1)$, and the words in R and B have length at most $2(n+1)$. Let \mathcal{P} be a set containing exactly one representative for each of these classes: \mathcal{P} is finite. Let

$$A = \max_{v \in V, P \in \mathcal{P}} \{|v|_P\}$$

Step 2: Choose a lift of I into \tilde{X} such that the induced injection of I in T_X contains z . This is possible since the foliation on X is arational, so the lifts of I in \tilde{X} intersect all the leaves. Let J be a closed interval $[a, b]$ of I whose interior contains z , and whose length is smaller than $\eta/2A$.

Each point of I is contained in at least two bands. Let D be the set of points contained in more than two bands: a point of D lies in a boundary component of Y , and arationality implies that it is the corner of exactly two bands. Conversely, if a point is the corner of a band, it is either contained in D , or it is one of the boundary p or q of I , which are also the corners of exactly two bands. Thus D has cardinal at most $2(n-1)$.

Let $d \in D$. There is a unique point d' of $\overset{\circ}{J}$ contained in the leaf of d , and such that the path joining d to d' in that leaf lies entirely in the interior of Y : indeed, if not, there would be a path in a leaf between two boundary components of Y , which contradicts arationality. Now let (r_a, r'_a) be the longest path in the leaf of a which contains a , lies entirely in the interior of Y , and does not meet $\overset{\circ}{J}$. The points r_a, r'_a are either points of J , or points of D . If r_a or r'_a is a point d of D , the leaf path $[r_a, r'_a]$ is contained in the leaf path $[d, d']$. We define similarly r_b, r'_b .

We now modify the band complex. For each point d of D , we cut Y along the path $[d, d']$. This enlarges the boundary component which contained d : if this component was the boundary of a 2-cell, we enlarge the 2-cell too. If both r_a and r'_a lie in $\overset{\circ}{J}$, we also cut along the leaf path $[r_a, r'_a]$, and glue a 2-cell along the boundary component thus created. We proceed similarly for b . It is straightforward to see that this gives us a new structure of band complex X' for Σ , whose canonical foliation is still homeomorphic to \mathcal{F} . Moreover, the union of bands Y' underlying X' is based in J .

Step 3: We want to see that the presentation associated to X' has at most $n+1$ generators and $2(n+1)$ relations of length at most $2(n+1)$. Suppose Y' is composed of r bands. This implies that the number of boundary components of Y' is at most $2r$ since each side of a band lies in exactly one boundary component. For the same reason, each boundary component is composed by at most $2r$ sides of bands. Now r is at most $n+1$: indeed, each point of D gives us a point d' contained in more than 2 bands of Y' , and both a and b might also give us such a point. But each such point is a corner of exactly two bands, and each band has at most four such points as corners. Since $|D| \leq 2(n-1)$, we get $r \leq n+1$, and Y' has at most $2(n+1)$ boundary components which are composed each of at most $2(n+1)$ sides of bands.

Step 4: This new structure of band complex gives us a presentation P_0 for G on at most $n+1$ generators, at most $2(n+1)$ relations, and in which the relations and boundary elements are represented by words of length at most $2(n+1)$ in the generators. Moreover, by remark 5.7, the generators of P_0 translate z by a distance less than the length of J , namely less than $\eta/2A$. By Step 1, there is a presentation P of our set of representatives \mathcal{P} which is combinatorially equivalent to P_0 . This gives us an automorphism $\phi_V : S \rightarrow S$ which sends any generator g of P on a generator of P_0 . Thus $d(z, \phi_V(g) \cdot z) < \eta/2A$, so that if $v \in V$, we have $d(z, \phi_V(v) \cdot z) < \eta$. This concludes the proof. \square

Recall that a vertex automorphism of a group G with respect to a splitting Λ is a standard extension of an automorphism of a vertex group of Λ to G .

The lemma we just proved is the key for dealing with paths of the form $[x, g \cdot x]$ which intersect at least one surface type component non-trivially. However, the argument needs to be completed, and caution is required, since a path can intersect several translates of a same surface type component.

The following result shows that if we know how to shorten simultaneously a finite number of paths that lie entirely in a surface component, then we can shorten paths which intersect non-trivially a translate of this component. This will enable us to prove Theorem 5.12 from lemma 5.13. In fact, it is more general, since it applies to any component in which orbits are dense: we will also use it to deal with axial components.

Lemma 5.15: *Let T be a G -tree which admits a decomposition as a graph of actions \mathcal{G} given by $(\tau, (T_v)_{v \in V(\tau)}, (p_e)_{e \in E(\tau)})$. Denote by $\Delta_{\mathcal{G}}$ the graph of groups corresponding to the action of G on τ . Let $G \cdot v$ be an orbit of vertices in τ such that*

1. *the action of G_v on T_v has dense orbits ;*
2. *for any finite set V of elements of G_v , for any point z of T_v and any positive η , there is an automorphism ϕ of G_v which restricts to a conjugation on each adjacent edge group, for which $d(z, \phi(g) \cdot z) < \eta$ for any element g in V .*

Then for any finite subset U of G , and for any $x \in T$, there exists a vertex automorphism τ of G relative to $\Delta_{\mathcal{G}}$ such that for $u \in U$, if $[x, u \cdot x]$ intersects a translate of T_v non-trivially, we have $d(x, \tau(u) \cdot x) < d(x, u \cdot x)$, and if not, $\tau(u) = u$.

The key to prove lemma 5.15 is to pick the right way of writing elements of G according to the splitting $\Delta_{\mathcal{G}}$, i.e. to choose the right isomorphism between G and $\pi_1(\Delta_{\mathcal{G}})$. This is precisely what the following lemma does.

Lemma 5.16: *Let T be a G -tree pointed by x which admits a decomposition as a graph of actions $\mathcal{G} = (\tau, (T_v)_{v \in V(\tau)}, (p_e)_{e \in E(\tau)})$. Denote by π the quotient map $\tau \rightarrow G \backslash \tau$, and by v_x the vertex of τ such that $x \in T_{v_x}$.*

Let v be a vertex in τ such that the action of G_v on T_v has dense orbits. Suppose that the path between v and v_x starts with an edge e_x , and does not meet any translates of v . Let z be the point of T_v closest to x . Let $\nu > 0$.

Let ϕ be an automorphism of G_v whose restriction to any adjacent edge group G_e is the conjugation by an element α_e of G_v , and assume that $\alpha_{e_x} = 1$.

There exists a fundamental domain τ_0 in τ containing the vertices v_x and v , and some Bass-Serre elements t_e in G for each edge e of τ_0 , such that for any e, f in τ_0

- *if $\pi(e)$ is adjacent to $\pi(v)$ then e is adjacent to v ;*
- *if $o(e) = o(f) = v$, and if p_e and p_f are in the same orbit, then $p_e = p_f$;*
- *if $o(e) = v$, we have $d(\alpha_e \cdot p_e, z) \leq \nu$ and if $\pi(t(e)) = \pi(v)$, we have $d(\alpha_e t_e^{-1} \cdot p_e, z) \leq \nu$;*
- *if $o(e) = v$ and p_e is in the same orbit as z , then $\alpha_e \cdot p_e = t_e^{-1} \cdot z$, and if $\pi(t(e)) = \pi(v)$ and p_e is in the same orbit as z , then $\alpha_e t_e^{-1} \cdot p_e = z$.*

Proof. We consider successively all the orbits of edges e of τ such that $o(e) = v$. For such an edge e , the point p_e lies in T_v . Note that if $e' = g \cdot e$ for $g \in G_v$, we have $\alpha_{e'} = \phi(g)\alpha_e g^{-1}$, and $p_{e'} = g \cdot p_e$.

If there is an element g of G_v such that $g \cdot p_e = p_f$ for p_f an edge that we already put in τ_0 , we choose the edge $e' = g \cdot e$ as a representative of the orbit of e in τ_0 . Then the edge group corresponding to e' is the stabiliser of $p_{e'} = p_f$, so it is the edge group corresponding to f . This implies $\alpha_{e'} = \alpha_f$, so that the condition $d(\alpha_{e'} \cdot p_{e'}, z) \leq \nu$ is satisfied.

Suppose now p_e is not in the orbit of any of the points p_f . Since the action of G_v on T_v has dense orbits, and since ϕ is an isomorphism, there is an element g of G_v such that $d(\phi(g)\alpha_e \cdot p_e, z) \leq \nu$. But $d(\phi(g)\alpha_e \cdot p_e, z) = d(\alpha_{e'} \cdot p_{e'}, z)$, so we choose the edge $g \cdot e$ as a representative of the orbit of e in τ_0 . If p_e is in the orbit of z , we take g such that $\phi(g)\alpha_e \cdot p_e = z$ so that $\alpha_{e'} \cdot p_{e'} = z$.

Denote by τ'_0 the connected subset of τ formed by all the edges we chose so far together with a representative w' for each orbit $G \cdot w$ where w is adjacent to v . Note that we have $\alpha_{e_x} \cdot p_{e_x} = z$ so we may assume that e_x lies in τ_0 . Thus we can extend τ'_0 to a fundamental domain τ_0 which contains v_x .

There remains to choose Bass-Serre elements for the edges of τ_0 . If e is an edge in τ_0 with $t(e) = g \cdot v$, then $g^{-1} \cdot p_e$ lies in T_v , so there is an element a of G_v such that $d(ag^{-1} \cdot p_e, z) \leq \nu$,

and $ag^{-1} \cdot p_e = z$ if $g^{-1} \cdot p_e$ and z are in the same orbit. We choose $t_e = \alpha_e^{-1}ga^{-1}$ as a Bass-Serre element for e . We extend this arbitrarily to a choice of Bass-Serre elements for all the edges of τ_0 . \square

Let us prove lemma 5.15.

Proof. Lemma 5.10 implies in particular that in our graph of actions, a path of the form $[x, g \cdot x]$ intersects finitely many components. Let thus ϵ be the minimal length, over all u in U , of the intersection of a path of the form $[x, u \cdot x]$ with a translate of T_v . Denote by v_x the vertex of τ such that $x \in T_{v_x}$. We may assume without loss of generality that the path between v_x and v does not go through any translates of v .

Given a fundamental domain τ_0 which contains v , and a choice of Bass-Serre elements $\{t_e\}_{e \in T_0}$ for the action of G on τ , denote by V the set of elements of G_v which appear in the loop representation of the elements u of U based at $\pi(v)$. The key remark is that this set is in fact independent of the choice of fundamental domain and Bass-Serre elements. By hypothesis, there is an automorphism ϕ of G_v such that $d(z, \phi(g) \cdot z) < \epsilon/8$ for any non-trivial element g of V . Moreover, the restriction of ϕ to each edge group G_e adjacent to G_v is a conjugation by some element α_e of G_v . Note that if e_x is the first edge of the path between v and v_x in τ , we may assume $\alpha_{e_x} = 1$.

Let us now choose a fundamental domain τ_0 which contains v , and some Bass-Serre elements $\{t_e\}_{e \in T_0}$ for the action of G on τ which satisfies the conclusions of 5.16 for $\nu = \epsilon/16$. This choice gives us an isomorphism between G and $\pi_1(\Delta_{\mathcal{G}})$, so we get a loop representaton based at $\pi(v)$ for any element u of G as

$$u = g_1 t_{e_1} g_2 t_{e_2} \dots g_l t_{e_l} g_{l+1}$$

Recall that the path formed by the edges $\pi(e_1), \dots, \pi(e_l)$ forms a loop based at $\pi(v_x)$, and that $g_i \in G_{o(\pi(e_i))}$. We can extend ϕ to G by taking a standard extension corresponding to the elements α_e and the choice of fundamental domain and Bass-Serre element we made. Note that then we have

$$\phi(u) = g'_1 t_{e_1} g'_2 t_{e_2} \dots g'_l t_{e_l} g'_{l+1},$$

where $g'_i = \alpha_{e_{i-1}}^{-1} \phi(g_i) \alpha_{e_i}$ if $g_i \in G_v$, and $g'_i = g_i$ if not.

According to lemma 5.10, the path $[x, u \cdot x]$ is the concatenation of the path $I_0 = [x, g_1 \cdot p_{e_1}]$, of translates I_1, \dots, I_{l-1} of paths of the form $[t_{e_i}^{-1} \cdot p_{e_i}, g_{i+1} \cdot p_{e_{i+1}}]$, and of a translate I_l of the path $[t_{e_l}^{-1} \cdot p_{e_l}, g_{l+1} \cdot x]$. Each of these paths lies in a different component of the graph of actions decomposition, we are interested in those that lie in translates of T_v . We have

$$d(x, u \cdot x) = \sum_{i=0}^l |I_i|. \quad (\dagger)$$

By the triangle inequality we have

$$d(x, \phi(u) \cdot x) \leq d(x, g'_1 \cdot p_{e_1}) + \sum_{i=1}^{l-1} d(t_{e_i}^{-1} \cdot p_{e_i}, g'_{i+1} \cdot p_{e_{i+1}}) + d(t_{e_l}^{-1} \cdot p_{e_l}, g'_{l+1} \cdot x)$$

We now want to compare this inequality to (\dagger) , for this we compare the summands of the right hand side to the lengths of the arcs I_i :

- If I_i is non-trivial and lies in a translate of T_v , we have $g_{i+1} \in V$, so we have $g'_{i+1} = \alpha_{e_i}^{-1} \phi(g_{i+1}) \alpha_{e_{i+1}}$ and $d(z, \phi(g_{i+1}) \cdot z) \leq \epsilon/8$. We get

$$\begin{aligned} d(t_{e_i}^{-1} \cdot p_{e_i}, g'_{i+1} \cdot p_{e_{i+1}}) &= d(\alpha_{e_i} t_{e_i}^{-1} \cdot p_{e_i}, \phi(g_{i+1}) \alpha_{e_{i+1}} \cdot p_{e_{i+1}}) \\ &\leq d(\alpha_{e_i} t_{e_i}^{-1} \cdot p_{e_i}, z) + d(z, \phi(g_{i+1}) \cdot z) + d(z, \alpha_{e_{i+1}} p_{e_{i+1}}) \leq \epsilon/4 \end{aligned}$$

But the length of I_i is at least ϵ , so we have

$$d(t_{e_i}^{-1} \cdot p_{e_i}, \psi(g_{i+1}) \cdot p_{e_{i+1}}) < d(t_{e_i}^{-1} \cdot p_{e_i}, g_{i+1} \cdot p_{e_{i+1}}) = |I_i|.$$

- If I_i does not lie in a translate of T_v , we have $g'_{i+1} = g_{i+1}$ so

$$d(t_{e_i}^{-1} \cdot p_{e_i}, g'_{i+1} \cdot p_{e_{i+1}}) = d(t_{e_i}^{-1} \cdot p_{e_i}, g_{i+1} \cdot p_{e_{i+1}}) = |I_i|$$

- If I_i trivial and lies in a translate of T_v , we have $t_{e_i}^{-1} \cdot p_{e_i} = g_{i+1} \cdot p_{e_{i+1}}$. But since $\pi(t(e_i)) = \pi(v)$, by the first point of lemma 5.16, $t_{e_i} = 1$ and in fact $p_{e_i} = g_{i+1} \cdot p_{e_{i+1}}$. We see that p_{e_i} and $p_{e_{i+1}}$ are in the same orbit, so by the second point our choice of fundamental domain, they are equal, so $\alpha_{e_i} = \alpha_{e_{i+1}}$ and we may assume $g_{i+1} = 1$. We get

$$\begin{aligned} d(t_{e_i}^{-1} \cdot p_{e_i}, g'_{i+1} \cdot p_{e_{i+1}}) &= d(t_{e_i}^{-1} \cdot p_{e_i}, \alpha_{e_i}^{-1} \phi(g_{i+1}) \alpha_{e_{i+1}} \cdot p_{e_{i+1}}) \\ &= d(t_{e_i}^{-1} \cdot p_{e_i}, p_{e_{i+1}}) \\ &= 0 = |I_i| \end{aligned}$$

Thus for $1 \leq i \leq l-1$, we see that $d(t_{e_i}^{-1} \cdot p_{e_i}, g'_{i+1} \cdot p_{e_{i+1}}) \leq |I_i|$ and the inequality is strict if and only if I_i is non-trivial and lies in a translate of T_v . Similarly, we can show that $d(x, g'_1 \cdot p_{e_1}) \leq |I_1|$ and that $d(t_{e_l}^{-1} \cdot p_{e_l}, g'_l \cdot x) \leq |I_{l+1}|$, and that these inequalities are strict if and only if I_1 and I_{l+1} respectively are non-trivial and lie in a translate of T_v . We get

$$d(x, \phi(u) \cdot x) \leq \sum_{i=1}^{l+1} |I_i| = d(x, u \cdot x)$$

and the inequality is strict if and only if at least one of the arcs I_i is non-trivial and lies in a translate of T_v . Thus if $[x, u \cdot x]$ intersects non-trivially a translate of T_v , it is made shorter by ϕ .

Now if $[x, u \cdot x]$ does not intersect any translates of T_v , for $1 \leq i \leq l+1$ we have $g'_i = g_i$ so $\phi(u) = u$. This finishes the proof. \square

We can now prove theorem 5.12.

Proof. Pick representatives T_1, \dots, T_r for the orbits of surface type components. Denote their stabilisers by S_1, \dots, S_r .

Lemma 5.13 tells us that the conditions of proposition 5.15 are satisfied. Thus, for each i , for any finite set of elements U_i of G , we can find ϕ_i such that if $u \in U_i$ and $[x, u \cdot x]$ intersects a translate of T_i non-trivially, we have $d(x, \phi_i(u) \cdot x) < d(x, u \cdot x)$, and if not, $\phi_i(u) = u$. Apply this successively to $U_1 = U, U_2 = \phi_1(U_1), \dots, U_r = \phi_{r-1}(U_{r-1})$, the automorphism $\phi = \phi_r \circ \dots \circ \phi_2 \circ \phi_1$ satisfies the conclusion of the theorem. \square

5.5 Axial Case

The theorem we need to deal with the axial case is very similar to that used for the surface case.

Theorem 5.17: *Let G be a finitely generated group whose solvable subgroups are free abelian. Suppose G acts on a real tree T with abelian pointwise arc stabilisers, and that T admits a decomposition as a graph of actions $\mathcal{G} = (\tau, (T_v)_{v \in V(\tau)}, (p_e)_{e \in E(\tau)})$. Denote by $\Delta_{\mathcal{G}}$ the graph of groups corresponding to the action of G on τ . Let U be a finite subset of G . There exists an element σ of $\text{Mod}(\Delta_{\mathcal{G}})$ such that for any element u of U ,*

- *if the geodesic segment $[x, u \cdot x]$ intersects non-trivially some axial components, then*

$$d(x, \sigma(u) \cdot x) < d(x, u \cdot x);$$

- *if not, $\sigma(u) = u$.*

We start by proving an analogue of 5.13.

Lemma 5.18: *Suppose that a finitely generated free abelian group A acts freely on a line L_0 by an action of axial type. Let V be a finite subset of A , let z be a point of T_0 and let $\eta > 0$. There exists an automorphism ϕ of A such that for any non-trivial element v of V ,*

$$d(z, \phi(v) \cdot z) < \eta.$$

Proof. Choose a basis a_1, \dots, a_k for A . Let K be the maximal length of an element of V with respect to this basis. Suppose without loss of generality that the elements a_i all translate L_0 in the same direction, and that the translation lengths are ordered as follows

$$\text{tr}(a_1) < \text{tr}(a_2) < \dots < \text{tr}(a_k)$$

Since the action is free, the translation lengths are \mathbb{Z} -independent. Thus for $i > 1$, there exists m_i such that $0 < \text{tr}(a_i) - m_i \text{tr}(a_1) < \text{tr}(a_1)$. So if ϕ is the isomorphism which sends a_i to $a_i - m_i a_1$ and fixes a_1 , we get

$$\text{tr}(\phi(a_i)) < \text{tr}(\phi(a_1)) = \text{tr}(a_1).$$

We can repeat this until the translation lengths of all the elements a_i are smaller than η/K . This proves the result. \square

Theorem 5.17 can now be proved.

Proof. Pick representatives T_1, \dots, T_r for the orbits of axial type components. Denote G_1, \dots, G_r their stabilisers. Each G_i is solvable, since it is an extension of a subgroup of $\text{Isom}(\mathbb{R})$, which is solvable, by a group fixing an arc of T , which is abelian by hypothesis. By hypothesis on G , we see that G_i is free abelian. We can thus write $G_i = A_0^i \oplus A_1^i$ where A_0^i acts trivially and A_1^i acts freely on T_i . By definition of an axial component, A_1^i is finitely generated.

Now if V is a finite subset of elements of G_i , let $V_1^i = p_{A_1^i}(V_i)$, the projection of V_i on A_1^i . By lemma 5.18, for any $\eta > 0$, and for any $z \in T_i$, there is an automorphism ϕ_i of A_1^i such that for all non-trivial element a of V_1^i , we have $d(z, \phi_i(a) \cdot z) \leq \eta$. We can extend ϕ_i to G_i by letting it be the identity on A_0^i : we get an automorphism of G_i which fixes all the edge groups adjacent to the vertex of τ corresponding to T_i , and for any element v of V which does not lie in A_0^i , we have $v = aw$ with $a \in A_1^i - \{1\}$ so $d(z, \phi_i(aw) \cdot z) = d(z, \phi_i(a) \cdot z) \leq \eta$.

Thus, by lemma 5.15, for each i , for any finite set of elements U_i of G , we can find ϕ_i such for any element u of U_i for which $[x, u \cdot x]$ intersects a translate of T_i non-trivially, we have $d(x, \phi_i(u) \cdot x) < d(x, u \cdot x)$, and if not, $\phi_i(u) = u$. If we apply this successively to $U_1 = U, U_2 = \phi_1(U_1), \dots, U_r = \phi_{r-1}(U_{r-1})$, the automorphism $\phi = \phi_r \circ \dots \circ \phi_2 \circ \phi_1$ satisfies the conclusion of the theorem. \square

5.6 Simplicial Case

We give the following definition.

Definition 5.19: (simplicial edge) *Suppose T is a real G -tree, denote by T' is minimal subtree. We call simplicial edge of T any non-trivial arc $[p, q]$ which lies in T' , whose interior contains no branching points in T' , and which is fixed pointwise by its stabiliser.*

Let $[p, q]$ be a simplicial edge in T with stabiliser C .

Definition 5.20: (splitting induced by a simplicial edge) *The group G admits a splitting Γ over C . We call this the splitting induced by $[p, q]$*

Denote by A and B the stabilisers of the connected components of $T' - \bigcup_{g \in G} g \cdot (p, q)$ which contain p and q respectively. Then either Γ is an amalgam of the form $A *_C B$, or there exists t in G such that $B = tAt^{-1}$ and Γ is an HNN extension of the form $A *_C$, with stable letter t .

Remark 5.21: If $[p, q]$ is a simplicial edge in T , any non-trivial subinterval of $[p, q]$ is also a simplicial edge of T , and it induces the same splitting as $[p, q]$. Indeed, since no branching points lie on $[p, q]$, the stabiliser of a non-trivial subarc is $\text{Stab}([p, q])$.

To take care of segments of the form $[x, u \cdot x]$ which lie entirely in the simplicial part, we will prove

Theorem 5.22: Let G be a finitely generated torsion-free group. Suppose $(X_n, x_n)_{n \in \mathbb{N}}$ is a sequence of δ_n -hyperbolic G -spaces which converges to a pointed real G -tree (T, x) whose arc stabilisers are abelian. Assume that in the stabiliser of any arc of T , there is an element which, for all n large enough, acts hyperbolically on X_n with translation length at least $12\delta_n$. Denote by T' the minimal subtree of T .

Let y be the point in T' of T which lies closest to x , and let U be a finite subset of elements of G . For all n large enough, there exists an automorphism ϕ_n of G , such that for any u in U ,

$$d(x, \phi_n(u) \cdot x) = d(x, u \cdot x);$$

and if the segment $[y, u \cdot y]$ contains a simplicial edge of T of the form $[y, q]$, we have

$$d_n(x_n, \phi_n(u) \cdot x_n) \leq d_n(x_n, u \cdot x_n)$$

with equality if and only if u fixes y .

Moreover, ϕ_n lies in the subgroup of $\text{Mod}(G)$ generated by the subgroups of the form $\text{Mod}(\Gamma)$, where Γ is the splitting induced by a simplicial edge of T .

To prove Theorem 5.22, we need to understand distances $d(x, g \cdot x)$ in the real G -tree T . Note first that $d(x, g \cdot x) = 2d(x, y) + d(y, g \cdot y)$ so we can restrict ourselves to understanding distances $d(y, g \cdot y)$ in the minimal tree.

Let $[y, q]$ be a simplicial edge in T , denote by Γ the splitting induced by $[y, q]$. Our aim now is to find a formula which expresses the distance $d(y, g \cdot y)$ in terms of the expression of g as a word given by the splitting Γ .

- If Γ is an amalgam $A *_C B$, we can write any element g of G as

$$g = a_1 b_1 \dots a_l b_l a_{l+1}$$

where $l \geq 0$ the elements a_i lie in A , the elements b_i lie in B , and they do not lie in C except possibly a_1 and a_{l+1} .

Then it is straightforward to show by induction that

$$d(y, g \cdot y) = \sum_{i=1}^{l+1} d(y, a_i \cdot y) + \sum_{i=1}^l d(y, b_i \cdot y). \quad (5.1)$$


Moreover, for any $b \in B - C$, we have $d(y, b \cdot y) = 2d(y, q) + d(q, b \cdot q)$ so

$$d(y, g \cdot y) = \sum_{i=1}^{l+1} d(y, a_i \cdot y) + 2ld(y, q) + \sum_{i=1}^l d(q, b_i \cdot q). \quad (5.2)$$

- If Γ is an HNN extension $A *_C$, we can write any element g of G as

$$g = a_0 t^{n_0} a_1 t^{n_1} a_2 \dots a_l t^{n_l} a_{l+1}$$

where the elements a_i are in A but not in C , except possibly a_0 and a_{l+1} , and $n_i \neq 0$ for $0 \leq i \leq l$. Let $r = t^{-1} \cdot q$.

Figure 5.1: HNN case: the path in T between y and $(a_1 t^{n_1} a_2 t^{n_2}) \cdot y$

We define for $0 \leq i \leq l$

$$\rho(i) = \begin{cases} y & \text{if } n_i \geq 0 \\ r & \text{if } n_i < 0 \end{cases} \quad \text{and} \quad \bar{\rho}(i) = \begin{cases} r & \text{if } n_i \geq 0 \\ y & \text{if } n_i < 0 \end{cases}$$

We also set $\rho(l+1) = y$.

It is fairly straightforward to see by induction (see figure 5.1) that the path between T_A and $a_j t^{n_j} a_{j+1} \dots a_l t^{n_l} a_{l+1} \cdot y$ starts at $a_j \cdot \rho(j)$ and that

$$d(y, g \cdot y) = d(y, a_0 \cdot \rho(0)) + d(T_A, t^{n_0} a_1 \dots a_l t^{n_l} \cdot y)$$

We deduce

$$d(y, g \cdot y) = d(y, a_0 \cdot \rho(0)) + \sum_{i=0}^l d(\rho(i), t^{n_i} \cdot \bar{\rho}(i)) + \sum_{i=0}^l d(\bar{\rho}(i), a_{i+1} \cdot \rho(i+1)) \quad (5.3)$$

Note also that $d(\rho(i), t^{n_i} \cdot \bar{\rho}(i)) = d(y, t^{|n_i|} \cdot r)$, and that for $k > 0$ we have

$$d(y, t^k \cdot r) = kd(y, q) + (k-1)d(r, y)$$

so if we let $N_g = \sum_{i=0}^l |n_i|$, we have

$$\begin{aligned} d(y, g \cdot y) &= d(y, a_0 \cdot \rho(0)) + N_g d(y, q) + (N_g - (l+1))d(r, y) \\ &\quad + \sum_{i=0}^l d(\bar{\rho}(i), a_{i+1} \cdot \rho(i+1)) \end{aligned} \quad (5.4)$$

With these result, we can show that the length of $[y, g \cdot y]$ in T does not change when we apply a Dehn twist of the splitting Γ to g .

Lemma 5.23: *Suppose T is a real G -tree with abelian arc stabilisers. Let $[y, q]$ be a simplicial edge in T , denote by Γ the splitting induced by $[y, q]$, and by A the stabiliser of the connected component of $T - \bigcup_{g \in G} g \cdot (y, q)$ which contains y . If ϕ is a Dehn twist by some element c of $\text{Stab}([y, q])$, for any element g of G we have*

$$d(y, g \cdot y) = d(y, \phi(g) \cdot y)$$

Proof. Suppose first that the splitting induced by $[y, q]$ is an amalgam $A *_C B$. Let $g \in G$, such that $g = a_1 b_1 \dots a_l b_l a_{l+1}$ where $l \geq 0$, the elements a_i lie in A , the elements b_i lie in B , and they do not lie in C except possibly a_1 and a_{l+1} .

Then we have by equation 5.1

$$d(y, g \cdot y) = \sum_{i=1}^{l+1} d(y, a_i \cdot y) + \sum_{i=1}^l d(y, b_i \cdot y).$$

With respect to the splitting Γ , the element $\phi(g)$ is represented by $a_1 (cb_1 c^{-1}) a_2 \dots a_l (cb_l c^{-1}) a_{l+1}$, so

$$d(y, \phi(g) \cdot y) = \sum_{i=1}^{l+1} d(y, a_i \cdot y) + \sum_{i=1}^l d(y, cb_i c^{-1} \cdot y).$$

But for any $b \in B$, we have $d(y, bcb^{-1} \cdot y) = d(c^{-1} \cdot y, bc^{-1} \cdot y) = d(y, b \cdot y)$ since c fixes y . Thus $d(y, \phi(g) \cdot y) = d(y, g \cdot y)$.

Consider now the case where the induced splitting is an HNN extension. Let g be an element of G , and choose a stable letter t . The element g can be written as $g = a_0 t^{n_0} a_1 t^{n_1} a_2 \dots a_l t^{n_l} a_{l+1}$ where the elements a_i are in A but not in C , except possibly for a_0 and a_{l+1} , and $n_i \neq 0$ for $0 \leq i \leq l$. By equation 5.3 we have

$$d(y, g \cdot y) = d(y, a_0 \cdot \rho(0)) + \sum_{i=0}^l d(\rho(i), t^{n_i} \cdot \bar{\rho}(i)) + \sum_{i=0}^l d(\bar{\rho}(i), a_{i+1} \cdot \rho(i+1))$$

Now $\phi(g) = a_0 (ct)^{n_0} a_1 (ct)^{n_1} a_2 \dots a_l (ct)^{n_l} a_{l+1}$, and this expression gives us a way to represent $\phi(g)$ in $A *_C$ with choice of stable letter $u = ct$. Thus equation 5.3 gives

$$d(y, \phi(g) \cdot y) = d(y, a_0 \cdot \rho(0)) + \sum_{i=0}^l d(\rho(i), (ct)^{n_i} \cdot \bar{\rho}(i)) + \sum_{i=0}^l d(\bar{\rho}(i), a_{i+1} \cdot \rho(i+1))$$

Now $d(y, (ct) \cdot r) = d(y, t \cdot r)$ since c fixes $q = t \cdot r$, so $d(y, \phi(g) \cdot y) = d(y, g \cdot y)$. \square

We are interested in segments of the form $[y, u \cdot y]$ which lie entirely in simplicial components. Such a segment must start with a simplicial edge of the form $[y, q]$. The following proposition enables us to shorten all paths of the form $[y, u \cdot y]$ which start with a given simplicial edge $[y, q]$. To prove Theorem 5.22, we will apply it to all simplicial edges of the form $[y, q]$.

Lemma 5.24: *Let G be a finitely generated torsion-free group. Suppose $(X_n, x_n)_{n \in \mathbb{N}}$ is a sequence of δ_n -hyperbolic G -spaces which converges to a pointed real G -tree (T, x) . Denote by T' the minimal subtree of T , and by y the point of T' closest to T' .*

Let $[y, q]$ be a simplicial edge of T whose stabiliser is non-trivial and abelian, and contains an element which, for all n large enough, acts hyperbolically on X_n with translation length at least $12\delta_n$. Denote by Γ the splitting induced by $[y, q]$, and by A the stabiliser of the connected components of $T - \bigcup_{g \in G} g \cdot (y, q)$ which contains y . Let V be a finite set of elements of G .

For any n large enough, there is a Dehn twist ϕ_n in $\text{Mod}(\Gamma)$ such that for any g in V ,

- *if $g \in A$, we have $\phi_n(g) = g$;*
- *if $g \notin A$, we have $d_n(x_n, \phi_n(g) \cdot x_n) < d_n(x_n, g \cdot x_n)$;*

To prove this we will use

Lemma 5.25: *Let G be a finitely generated torsion-free group. Suppose $(X_n, x_n)_{n \in \mathbb{N}}$ is a sequence of δ_n -hyperbolic G -spaces which converges to a pointed real G -tree (T, x) . Let $[y, q]$ be a simplicial edge of T . Suppose that there exists a non-trivial element c fixing $[y, q]$, such that for n large enough, c acts hyperbolically on X_n , with translation length at least $12\delta_n$. Denote by Γ the splitting induced by $[y, q]$.*

For any $\epsilon > 0$, for n large enough, if q_n and y_n approximate q and y in an ϵ -approximation between T and X_n with respect to y, q and c , there exists an integer k_n such that

$$d_n(y_n, c^{k_n} \cdot q_n) \leq 5\epsilon.$$

Proof. Let $\epsilon > 0$. Fix n large enough so that c acts hyperbolically on X_n with translation length at least $12\delta_n$, and such that $12\delta_n < \epsilon/3$. Suppose that there exists an ϵ -approximation between T and X_n with respect to y, q and c . Denote by q_n and y_n the approximations of q and y respectively.

Choose a geodesic L contained in $\text{Ax}(c)$ (recall $\text{Ax}(c)$ is the set of geodesics joining the points of the boundary fixed by c), and a geodesic parametrisation $t \mapsto w(t)$ of L .

In T , the element c fixes y and q , so $d_n(y_n, c \cdot y_n) < \epsilon$ and $d_n(q_n, c \cdot q_n) < \epsilon$. By lemma 3.13, the axis of c is at a distance at most $4\delta_n$ of the midpoints of $[y_n, c \cdot y_n]$ and $[q_n, c \cdot q_n]$. Thus these midpoints are at a distance at most $12\delta_n$ of L , so that $d_n(y_n, L) < \epsilon$ and $d_n(q_n, L) < \epsilon$. Without loss of generality, we may assume that $d_n(q_n, w(0)) < \epsilon$, and that there is a positive real u such that $d_n(y_n, w(u)) < \epsilon$.

The elements of the form c^k also acts hyperbolically on X_n , and their axis coincides with the axis of c since they fix the same points on the boundary. Let us consider their translation lengths $\text{tr}_n(c^k)$. For any point z of $\text{Ax}_n(c)$, we have $d_n(z, c^k \cdot z) \leq \text{tr}_n(c^k) + 16\delta_n$ by lemma 3.11. But the map $\mathbb{Z} \rightarrow X$ given by $k \mapsto c^k \cdot z$ is a quasi-isometry, so $\text{tr}_n(c^k)$ must tend to ∞ as k tends to infinity.

On the other hand, $d_n(w(u), c \cdot w(u)) \leq 2d_n(y_n, w(u)) + d_n(y_n, c \cdot y_n) \leq 3\epsilon$. Thus for $k \in \mathbb{Z}$, we have

$$\begin{aligned} |\text{tr}_n(c^{k+1}) - \text{tr}_n(c^k)| &\leq |d_n(w(u), c^{k+1} \cdot w(u)) - d(w(u), c^k \cdot w(u))| + 32\delta_n \\ &\leq |d_n(c^k \cdot w(u), c^{k+1} \cdot w(u))| + 32\delta_n \\ &< 4\epsilon. \end{aligned}$$

The sequence $(\text{tr}_n(c^k))_{k \in \mathbb{N}}$ tends to infinity, and the distance between two consecutive terms is at most 4ϵ , so there exists a positive integer k_n such that $|u - \text{tr}_n(c^{k_n})| < 2\epsilon$.

The element c^{k_n} acts as a $20\delta_n$ -quasitranslation of length $\text{tr}(c^{k_n})$ on the geodesic $t \mapsto w(t)$, and we get

$$d_n(w(\text{tr}_n(c^{k_n})), c^{k_n} \cdot w(0)) < 20\delta_n < \epsilon.$$

But $d_n(w(u), w(\text{tr}_n(c^{k_n}))) = |u - \text{tr}_n(c^{k_n})| < 2\epsilon$ and we deduce

$$\begin{aligned} d_n(y_n, c^{k_n} \cdot q_n) &\leq d_n(y_n, w(u)) + d_n(w(u), c^{k_n} \cdot w(0)) + d_n(w(0), q_n) \\ &< d_n(w(u), c^{k_n} \cdot w(0)) + 2\epsilon \\ &\leq d_n(w(u), w(\text{tr}_n(c^{k_n}))) + d_n(w(\text{tr}_n(c^{k_n})), c^{k_n} \cdot w(0)) + 2\epsilon \\ &\leq 5\epsilon \end{aligned}$$

This terminates the proof. □

We can now prove lemma 5.24.

Proof. We consider separately the case where Γ is an amalgam, and the case where it is an HNN extension.

Amalgam case: Write each element g of V according to this splitting as $g = a_1 b_1 \dots a_l b_l a_{l+1}$, where $l \geq 1$, the elements a_i lie in A , the elements b_i lie in B , and they do not lie in C except possibly a_1 and a_{l+1} .

Let V_A and V_B be the finite sets of elements a_i and b_i respectively that appear in all such decompositions. Let ϵ be such that $0 < \epsilon < d(y, q)/100$. Let c be the non-trivial element of C given by the hypotheses.

Let n be large enough so that there exists an ϵ -approximation R between T and X_n relative to y, q and $V \cup V_A \cup V_B \cup \{c\}$. By lemma 5.25, if n is large enough, there exists an integer k_n such that $d_n(y_n, c^{k_n} \cdot q_n) < 5\epsilon$.

Let g be an element of V which does not lie in A . It admits a decomposition according to the splitting Γ given by $g = a_1 b_1 \dots a_l b_l a_{l+1}$, and $l > 0$ since $g \notin A$. By equation 5.2, we have

$$d(y, g \cdot y) = \sum_{i=1}^{l+1} d(y, a_i \cdot y) + \sum_{i=1}^l d(q, b_i \cdot q) + 2ld(y, q)$$

so by our ϵ approximation, we get

$$d_n(y_n, g \cdot y_n) \geq \sum_{i=1}^{l+1} d_n(y_n, a_i \cdot y_n) + \sum_{i=1}^l d_n(q_n, b_i \cdot q_n) + 2ld_n(y_n, q_n) - (4l + 2)\epsilon \quad (5.5)$$

Let ϕ_n be the Dehn twist about c^{k_n} . The triangle inequality gives

$$d_n(y_n, \phi_n(g) \cdot y_n) \leq \sum_{i=1}^{l+1} d_n(y_n, a_i \cdot y_n) + \sum_{i=1}^l d_n(y_n, c^{k_n} b_i c^{-k_n} \cdot y_n)$$

But for $b \in B - C$, we have by the triangle inequality

$$\begin{aligned} d_n(y_n, c^{k_n} b c^{-k_n} \cdot y_n) &\leq d_n(y_n, c^{k_n} \cdot q_n) + d_n(c^{k_n} \cdot q_n, c^{k_n} b \cdot q_n) + d_n(c^{k_n} b \cdot q_n, c^{k_n} b c^{-k_n} \cdot y_n) \\ &= 10\epsilon + d_n(q_n, b \cdot q_n). \end{aligned}$$

We finally get

$$\begin{aligned} d_n(y_n, \phi_n(g) \cdot y_n) &\leq \sum_{i=1}^{l+1} d_n(y_n, a_i \cdot y_n) + \sum_{i=1}^l d_n(q_n, b_i \cdot q_n) + 10l\epsilon \\ &\leq d_n(y_n, g \cdot y_n) - 2ld_n(y_n, q_n) + (14l + 2)\epsilon \end{aligned}$$

by equation 5.5. Now $d(x, \phi_n(g) \cdot x) = 2d(x, y) + d(y, \phi_n(g) \cdot y)$, so $d_n(x_n, \phi(g) \cdot x_n) > 2d_n(x_n, y_n) + d_n(y_n, g \cdot y_n) - 4\epsilon$. We get

$$\begin{aligned} d_n(x_n, \phi_n(g) \cdot x_n) &\leq 2d_n(x_n, y_n) + d_n(y_n, \phi_n(g) \cdot y_n) \\ &\leq 2d_n(x_n, y_n) + d_n(y_n, g \cdot y_n) - 2ld_n(y_n, q_n) + (14l + 2)\epsilon \\ &\leq d_n(x_n, g \cdot x_n) - 2ld_n(y_n, q_n) + (14l + 6)\epsilon \\ &< d_n(x_n, g \cdot x_n) \end{aligned}$$

since $\epsilon < d(y, q)/100 < d_n(y_n, q_n)/99$ so $(14l + 6)\epsilon < 2ld_n(y_n, q_n)$ (recall that $l > 0$).

HNN case: Let us now consider the case where Γ is an HNN extension. We choose a stable letter t such that $q \in t \cdot T_A$, and we write each element g of V as $a_0 t^{n_0} a_1 t^{n_1} a_2 \dots a_l t^{n_l} a_{l+1}$, where the elements a_i are in A but not in C , except possibly a_0 and a_{l+1} , and $n_i \neq 0$ for $0 < i < l$. Denote by N_g the integer $\sum_{i=0}^l |n_i|$. Note that if g is not in A , we must have $N_g > 0$.

Let V_A be the finite set of elements a_i that appear in all such decompositions for $g \notin A$. Let ϵ be such that $0 < \epsilon < d(y, q)/100$. Let n be large enough so that there exists an ϵ -approximation R between T and X_n relative to $y, q, r = t^{-1} \cdot q$ and $V \cup V_A \cup \{t\}$. There exists an element c of C such that for all n large enough, $\text{tr}(c) > 12\delta_n$. Thus by lemma 5.25, for n large enough there is an integer k_n such that $d_n(y_n, c^{k_n} \cdot q_n) < 5\epsilon$, where q_n is a point approximating q with respect to R . Thus we have $d_n(y_n, c^{k_n} t \cdot r_n) < 6\epsilon$.

Let g be an element of V , whose decomposition according to the splitting Γ is given by $a_0 t^{n_0} a_1 t^{n_1} a_2 \dots a_l t^{n_l} a_{l+1}$. By equation 5.4, we have

$$d(y, g \cdot y) = d(y, a_0 \cdot \rho(0)) + \sum_{i=0}^l d(\bar{\rho}(i), a_{i+1} \cdot \rho(i+1)) + (N_g - l - 1)d(r, y) + N_g d(y, q)$$

so by our ϵ -approximation we get

$$\begin{aligned} d_n(y_n, g \cdot y_n) &\geq d_n(y_n, a_0 \cdot \rho_n(0)) + \sum_{i=0}^l d_n(\bar{\rho}_n(i), a_{i+1} \cdot \rho_n(i+1)) \\ &\quad + (N_g - l - 1)d_n(r_n, y_n) + N_g d_n(y_n, q_n) - (2N_g + 2)\epsilon \end{aligned} \quad (5.6)$$

Let ϕ_n be the Dehn twist about c^{k_n} . We have by the triangle inequality

$$\begin{aligned} d_n(y_n, \phi_n(g) \cdot y_n) &\leq d_n(y_n, a_0 \cdot \rho_n(0)) + \sum_{i=0}^l d_n(\bar{\rho}_n(i), a_{i+1} \cdot \rho_n(i+1)) \\ &\quad + \sum_{i=0}^l d_n(\rho_n(i), (c^{k_n} t)^{n_i} \cdot \bar{\rho}_n(i)) \end{aligned}$$

Note that $d_n(\rho_n(i), (c^{k_n} t)^{n_i} \cdot \bar{\rho}_n(i)) = d_n(y_n, (c^{k_n} t)^{|n_i|} \cdot r_n)$. By the triangle inequality, we have for $j > 0$

$$\begin{aligned} d_n(y_n, (c^{k_n} t)^j \cdot r_n) &\leq d_n(y_n, c^{k_n} t \cdot r_n) + d_n(c^{k_n} t \cdot r_n, (c^{k_n} t)^j \cdot r_n) \\ &\leq d_n(y_n, c^{k_n} t \cdot r_n) + (j-1)d_n(r_n, c^{k_n} t \cdot r_n) \\ &\leq j d_n(y_n, c^{k_n} t \cdot r_n) + (j-1)d_n(r_n, y_n) \\ &\leq (j-1)d_n(r_n, y_n) + 6j\epsilon. \end{aligned}$$

Thus we see that

$$\begin{aligned} d_n(y_n, \phi_n(g) \cdot y_n) &\leq d_n(y_n, a_0 \cdot \rho_n(0)) + \sum_{i=0}^l d_n(\bar{\rho}_n(i), a_{i+1} \cdot \rho_n(i+1)) \\ &\quad + (N_g - l - 1)d_n(r_n, y_n) + 6N_g\epsilon \\ &\leq d_n(y_n, g \cdot y_n) - N_g d_n(y_n, q_n) + (8N_g + 2)\epsilon \end{aligned}$$

by equation 5.6. Now $d(x, \phi_n(g) \cdot x) = 2d(x, y) + d(y, \phi_n(g) \cdot y)$, so $d_n(x_n, \phi(g) \cdot x_n) > 2d_n(x_n, y_n) + d_n(y_n, g \cdot y_n) - 4\epsilon$. We get

$$d_n(x_n, \phi_n(g) \cdot x_n) \leq 2d_n(x_n, y_n) + d_n(y_n, \phi_n(g) \cdot y_n) \quad (5.7)$$

$$\leq 2d_n(x_n, y_n) + d_n(y_n, g \cdot y_n) - N_g d_n(y_n, q_n) + (8N_g + 2)\epsilon \quad (5.8)$$

$$\leq d_n(x_n, g \cdot x_n) - N_g d_n(y_n, q_n) + (8N_g + 6)\epsilon \quad (5.9)$$

$$< d_n(x_n, g \cdot x_n) \quad (5.10)$$

since $\epsilon < d(y, q)/100 < d_n(y_n, q_n)/99$ so $(8N_g + 6)\epsilon < N_g d_n(y_n, q_n)$ (recall $N_g > 0$ since $g \notin A$). \square

We can now prove Theorem 5.22.

Proof. There is a finite number of orbits of maximal (for inclusion) simplicial edges of the form $[y, q]$ in T : we choose some representatives $[y, q_1], \dots, [y, q_m]$. These edges induce splittings $\Gamma_1, \dots, \Gamma_m$. Each splitting Γ_i is of the form $A_i *_{C_i} B_i$ or $A_i *_ {C_i}$, where A_i stabilises the connected component of $T - \bigcup_{g \in G} g \cdot (y, q_i)$ which contains y .

If $g \in (A_1 \cap \dots \cap A_m)$, we see that the path $[y, g \cdot y]$ does not intersect any translates of the edges $[y, q_i]$. By lemma 5.24, for all n large enough, we can find a Dehn twist ϕ_n^1 in $\text{Mod}(\Gamma_1)$ such that for any $g \in U$;

- if $g \in A_1$, we have $\phi_n^1(g) = g$;
- if $g \notin A_1$, we have $d_n(x_n, \phi_n^1(g) \cdot x_n) < d_n(x_n, g \cdot x_n)$.

We apply lemma 5.24 successively to the sets $V = U \cap (A_1 \cap \dots \cap A_{i-1})$, to find for any n large enough a Dehn twist ϕ_n^i in $\text{Mod}(\Gamma_i)$ such that for any $g \in U$;

- if $g \in (A_1 \cap \dots \cap A_i)$, we have $(\phi_n^i \circ \dots \circ \phi_n^1)(g) = g$;
- if $g \notin (A_1 \cap \dots \cap A_i)$, we have $d_n(x_n, (\phi_n^i \circ \dots \circ \phi_n^1)(g) \cdot x_n) < d_n(x_n, g \cdot x_n)$.

Finally we set $\phi_n = \phi_n^m \circ \dots \circ \phi_n^1$. If $g \in U$ is such that $[y, g \cdot y]$ contains a simplicial edge of the form $[y, q]$, we know that $g \notin (A_1 \cap \dots \cap A_m)$, so $d_n(x_n, \phi_n(g) \cdot x_n) < d_n(x_n, g \cdot x_n)$.

Since all the automorphisms ϕ_i are Dehn twists, by lemma 5.23, we see that for any $g \in G$ we have $d(y, \phi_n(g) \cdot y) = d(y, g \cdot y)$ so

$$d(x, \phi_n(g) \cdot x) = 2d(x, y) + d(y, \phi_n(g) \cdot y) = 2d(x, y) + d(y, g \cdot y) = d(x, g \cdot x).$$

□

5.7 Proof of the shortening Theorem

Putting all of the pieces together, we can now prove Theorem 4.28.

Proof. Let T' be the minimal subtree of the action of G on T . Let y be the point of T' such that $[x, y]$ is the shortest path between x and T' . Note that H fixes a point in T' .

If g is an element of G , we have

$$d(x, g \cdot x) = 2d(x, y) + d(y, g \cdot y).$$

The action of g on T is non-trivial, so there is at least one elements g_0 in the generating set Σ_G , such that $g_0 \cdot y \neq y$. Moreover, if an element g of Σ_G fixes y , the distance $d(x, g \cdot x)$ is strictly smaller than $d(x, g_0 \cdot x)$, so that for n large enough, $d_n(x_n, g \cdot x_n)$ is strictly smaller than $d_n(x_n, g_0 \cdot x_n)$. The maximal displacement of the basepoint is reached by a generator which does not fix y . Therefore, we can assume that none of the elements of Σ_G fix y .

The tree (T', y) satisfies the hypotheses of Theorem 5.11, so it admits a graph of actions $\mathcal{G} = (\tau, (T_v)_{v \in V(\tau)}, (p_e)_{e \in E(\tau)})$ into surface, axial and simplicial components. Denote by $\Delta_{\mathcal{G}}$ the graph of groups corresponding to the action of G on τ . Denote by $\Gamma_{\mathcal{G}}$ the refinement of $\Delta_{\mathcal{G}}$ by the actions of the simplicial type vertices on their simplicial vertex trees. Note that $\text{Mod}(\Delta_{\mathcal{G}}) < \text{Mod}(\Gamma_{\mathcal{G}})$.

We apply proposition 5.12 to the set $U = \Sigma_G$, to get an element ϕ_s of $\text{Mod}(\Delta_{\mathcal{G}})$ such that for every g in Σ_G , we have $d(y, \phi_s(g) \cdot y) < d(y, g \cdot y)$, unless $[y, g \cdot y]$ does not intersect any surface type components, in which case $\phi_s(g) = g$.

We then apply proposition 5.17 to the set $U' = \phi_s(\Sigma_G)$, to get an element ϕ_a of $\text{Mod}(\Delta_G)$ such that for every g in Σ_G , we have $d(y, \phi_a(\phi_s(g)) \cdot y) < d(y, \phi_s(g) \cdot y)$, unless $[y, \phi_s(g) \cdot y]$ does not intersect any axial type components, in which case $\phi_a(\phi_s(g)) = \phi_s(g)$.

Thus

$$\begin{aligned} d(x, \phi_a(\phi_s(g)) \cdot x) &= 2d(x, y) + d(y, \phi_a(\phi_s(g)) \cdot y) \\ &< 2d(x, y) + d(y, g \cdot y) \\ &= d(x, g \cdot x) \end{aligned}$$

unless $[y, g \cdot y]$ lies entirely in the simplicial part of T' , in which case $\phi_a(\phi_s(g)) = g$. Let $\delta > 0$ be smaller than $d(x, g \cdot x) - d(x, \phi_a(\phi_s(g)) \cdot x)$ for all the elements g of Σ_G such that $[y, g \cdot y]$ intersects a surface or an axial component.

Let n be large enough so that there is a $\delta/2$ -approximation between X_n and T relative to $\phi_a(\phi_s(\Sigma_G))$, and so that proposition 5.22 applied to the set $U'' = \phi_a \circ \phi_s(\Sigma_G)$ gives us an element ϕ_n of $\text{Mod}(\Gamma_G)$ for which

- $d(x, \phi_n \circ \phi_a \circ \phi_s(g) \cdot x) = d(x, \phi_a \circ \phi_s(g) \cdot x)$ for all g in Σ_G ;
- $d_n(x_n, \phi_n \circ \phi_a \circ \phi_s(g) \cdot x_n) < d_n(x_n, \phi_a \circ \phi_s(g) \cdot x_n)$ for all g such that $[y, \phi_a \circ \phi_s(g) \cdot y]$ contains a simplicial edge of the form $[y, q]$.

Let $\phi = \phi_n \circ \phi_a \circ \phi_s$. Note that since H fixes a point in T' , the modular group $\text{Mod}(\Gamma_G)$ is a subgroup of $\text{Mod}_H(G)$. For g in Σ_G such that $[y, g \cdot y]$ intersects a surface or an axial component, we have:

$$\begin{aligned} d_n(x_n, \phi(g) \cdot x_n) &< d(x, \phi(g) \cdot x) + \delta/2 \\ &= d(x, \phi_a \circ \phi_s(g) \cdot x) + \delta/2 \\ &\leq d(x, g \cdot x) - \delta/2 \\ &\leq d_n(x_n, g \cdot x_n). \end{aligned}$$

If g is an element of Σ_G such that $[y, g \cdot y]$ lies entirely in the simplicial part of T' , we have $\phi_a \circ \phi_s(g) = g$, so $[y, \phi_a \circ \phi_s(g) \cdot y]$ lies entirely in the simplicial part of T' , hence must start with a simplicial edge of T . Therefore $d_n(x_n, \phi(g) \cdot x_n) = d_n(x_n, \phi_n \circ \phi_a \circ \phi_s(g) \cdot x_n) < d_n(x_n, g \cdot x_n)$. This finishes the proof. \square

Chapter 6

Factor sets

The results presented in this section form an essential step of the construction of a Makanin-Razborov diagram, which analyses the set of morphisms $\text{Hom}(G, \Gamma)$ from a given finitely generated group G into a free group or into a hyperbolic group Γ for example. We prove the existence of a factor set for such morphisms, that is, we show that there is a finite number of morphisms f_1, \dots, f_m such that, up to precomposition by an automorphism, any element of $\text{Hom}(G, \Gamma)$ factors through one of the maps f_i .

6.1 Case of free groups

Recall that a sequence of morphisms $h_n : G \rightarrow G'$ is stable if for any element g of G , either $h_n(g)$ is trivial for all but finitely many values of n , or $h_n(g)$ is non-trivial for all but finitely many values of n . Recall also that the set of elements g whose image by h_n is almost everywhere trivial is called the stable kernel of the sequence, and denoted by $\underline{\text{Ker}}(h_n)$.

Definition 6.1: (limit of a stable sequence) *The limit of a stable sequence $h_n : G \rightarrow G'$ is the group $G/\underline{\text{Ker}}(h_n)$.*

Definition 6.2: (limit group) *A limit group L is the limit of a stable sequence of morphisms from a finitely generated group G into a free group $f_n : G \rightarrow \mathbb{F}$.*

The following proposition lists some properties of limit groups that will be of use. All these are elementary, and proved in lemma 1.4 of [Sel01] or proposition 3.1 of [CG05].

Proposition 6.3: *Let L be a limit group.*

- *L is torsion-free;*
- *maximal abelian subgroups of L are malnormal;*
- *a solvable subgroup of L is abelian;*
- *given two elements a, b in L either a and b commute, or they generate a free group of rank 2.*

We will also use the following property, proved in [Sel01].

Proposition 6.4: *Limit groups are finitely presented.*

This is a highly non-trivial fact. To show finite presentability of limit groups, Sela shows that the class of limit groups coincides with that of constructible limit groups (see section 4 of [Sel01],

alternatively this proof is also written up in [Wil06] and [BF03]). Constructible limit groups are easily seen to be finitely presented, hence the result.

It is also fairly straightforward to see that constructible groups have finitely generated abelian subgroups, thus as a corollary we get

Proposition 6.5: *Abelian subgroups of a limit group are finitely generated.*

From this we can deduce in particular

Lemma 6.6: *A virtually solvable subgroup of a limit group is free abelian.*

Proof. We know by the third point of 6.3 that solvable subgroups of limit groups are in fact abelian. We just saw that abelian subgroups of limit groups are finitely generated. Thus virtually solvable subgroups of a limit group are finitely generated, and have polynomial growth. This implies in particular that a virtually solvable subgroup H cannot contain a free group of rank 2, so by the fourth point of 6.3, any two elements of H must commute. Thus H is finitely generated abelian, and since L is torsion free, it must in fact be free abelian. \square

Definition 6.7: (shortening quotient) *A shortening quotient of a group G is the limit of a stable sequence of morphisms $h_n : G \rightarrow \mathbb{F}$ which are short in the sense of definition 4.22.*

Theorem 4.25 says that if G satisfies some nice properties, and if the h_n are short and non-injective, the limit group L is a proper quotient. Note that if G is both freely indecomposable and non-cyclic, it cannot inject into a free group, so the non-injectivity of the h_n is automatically satisfied. The following can thus be seen as yet another version of the shortening argument.

Theorem 6.8: *Suppose G is a freely indecomposable, torsion-free and non-cyclic finitely generated group. Suppose moreover that its virtually solvable subgroups are free abelian. Then a shortening quotient of G is a proper quotient.*

We want to show the following result:

Proposition 6.9: *Let G be a non-cyclic and freely indecomposable finitely generated group. There exist a finite set of limit groups which are proper quotients of G such that any morphism f from G to a free group \mathbb{F} factors through one of the corresponding quotient maps after precomposition by a modular automorphism.*

Definition 6.10: (factor set) *Such a finite set of quotients is called a factor set for $\text{Hom}(G, \mathbb{F})$.*

To do this, we introduce a partial order relation on the shortening quotients of G : suppose L_1, L_2 are shortening quotients with quotient maps $\eta_i : G \rightarrow L_i$, for $i = 1, 2$. We say that $L_1 \leq L_2$ if the map η_1 factors through η_2 , i.e. if there exists a map $\nu : L_2 \rightarrow L_1$ such that $\eta_1 = \nu \circ \eta_2$. This amounts to saying that $\text{Ker}(\eta_2) \subseteq \text{Ker}(\eta_1)$.

We now show that every shortening quotient is smaller than a maximal shortening quotient for this order relation, and that there is only a finite number of maximal shortening quotient.

Proposition 6.11: *Let G be a non-cyclic and freely indecomposable finitely generated group. Every shortening quotient of G is smaller than a maximal shortening quotient.*

Proof. We will of course apply Zorn's lemma, for this we need to show that every totally ordered set of shortening quotients has an upper bound.

Assume without loss of generality that the totally ordered set of shortening quotients is infinite. It contains a strictly increasing sequence $Q_1 < Q_2 < Q_3 \dots$ of shortening quotients of G , with corresponding maps $\eta_n : G \rightarrow Q_n$, such that η_n minimises the set $\text{Ker}(\eta) \cap B_G(n)$. The kernel of any quotient $\eta : G \rightarrow Q$ in our totally ordered set contains $\text{Ker}(\eta) \cap B_G(n)$. Suppose we can find an upper bound $\eta_\infty : G \rightarrow Q_\infty$ for this sequence. If g is an element of $\text{Ker}(\eta_\infty)$, then $g \in \text{Ker}(\eta_n)$ for all indices n , thus $g \in \text{Ker}(\eta_{l(g)})$. But for any quotient $\eta : G \rightarrow Q$ in our ordered set,

$\text{Ker}(\eta_{(g)}) \cap B_G(l(g))$ is contained $\text{Ker}(\eta)$. Thus it is enough to show that any ordered sequence of shortening quotient has an upper bound.

Each Q_n is the limit of a sequence of short morphisms $h_i^n : G \rightarrow \mathbb{F}$ (since G is finitely generated we can assume the group \mathbb{F} is the same for all n). By extracting carefully from each sequence $(h_i^n)_i$, we can moreover assume that for each $i \in \mathbb{N}$, for any word w in G of length at most i , $h_i^n(w) = 1$ if and only if $\eta_n(w) = 1$. Consider the diagonal sequence of short morphisms $h_j^j : G \rightarrow \mathbb{F}$. Extract a stable subsequence (which we still denote h_j^j). Denote by Q the limit group which is the limit of this sequence, and by η the quotient map. It is a shortening quotient since all the maps h_j^j are short.

Now Q , as a limit group, is finitely presented. Thus for j large enough, all the elements of the kernel of η are mapped to 1 by h_j^j , that is, if $\eta(w) = 1$, then $\eta_j(w) = h_j^j(w) = 1$. Thus $Q_j < Q$ for j large enough, so this holds for all j . This terminates the proof. \square

Proposition 6.12: *There is only a finite number of maximal shortening quotient.*

Proof. Assume $(M_n)_{n \in \mathbb{N}}$ is an infinite sequence of maximal shortening quotients of G . Each M_n is the limit of a sequence $(h_i^n)_{i \in \mathbb{N}}$ of short morphisms $G \rightarrow \mathbb{F}$ which we choose again to ensure that the kernels of h_i^n and η_n coincide on words of length less than or equal to i . Extract from the diagonal sequence of morphisms h_j^j a stable subsequence (still denoted h_j^j). Let M be the shortening quotient of G limit of this sequence, η the corresponding quotient map.

As M is finitely presented, and the sequence h_j^j is stable, for j large enough h_j^j maps all the elements of the kernel of η to 1. As h_j^j agrees with η_j on words of length less than or equal to j , for j large enough η_j sends all the elements of the kernel of η to 1, in other words M_j is a quotient of M . But M_j is maximal, so M is equivalent to M_j . Since this holds for all j large enough, the quotients M_j cannot be all distinct. \square

We can now prove proposition 6.9:

Proof. By 6.12, G has a finite set M_1, \dots, M_r of maximal shortening quotients of G with corresponding quotient maps η_1, \dots, η_r .

If G is not a limit group, a shortening quotient of G must be a proper quotient since a shortening quotient is in particular a limit group.

If G is a limit group, it is torsion-free and its virtually solvable subgroups are free abelian by 6.6. Thus by proposition 6.8, its shortening quotients are proper quotients

In both cases the surjective maps η_1, \dots, η_r are not injective. Let f be a morphism $G \rightarrow \mathbb{F}$. Let $\sigma \in \text{Mod}(G)$ be such that $\bar{f} = f \circ \sigma$ is short. The limit of the sequence \bar{f}, \bar{f}, \dots is a shortening quotient Q of G , with quotient map η . Note that \bar{f} factors through η . By proposition 6.11, η factors through one of the η_j . But this implies \bar{f} factors through η_j , thus proving the claim. \square

6.2 Case of torsion-free hyperbolic groups

Let Γ be a torsion-free hyperbolic group finitely generated by a set Σ .

Definition 6.13: (Γ -limit group, strict Γ -limit group) *A Γ -limit group L is the limit of a stable sequence of morphisms from a finitely generated group G into Γ . If moreover the morphisms are pairwise non-conjugate, we say that L is a strict Γ -limit group.*

Remark 6.14: *A Γ -limit group L is either a strict Γ -limit group or a finitely generated subgroup of Γ . Indeed, from the sequence of morphisms which defines it, we can extract either a subsequence of non-conjugate morphisms, or a subsequence of pairwise conjugate morphisms, and L is still*

the limit of this subsequence. In the first case L is a strict Γ -limit group, in the second case L is isomorphic to the subgroup of Γ image of any of the morphisms.

At this point, what we did in the free case was to state the result that limit groups are finitely presented. This is not necessarily true in the case of Γ -limit group, indeed Γ may have finitely generated subgroups that are not finitely presented.

However it is true that a chain of strict epimorphisms between Γ -limit groups stabilises. Note that this would be obvious if we knew them to be finitely presented. The proof requires heavy construction, in particular the 'shortening procedure', and can be found in Sela and as Theorem 5.2 in [Gro05] (here Groves treats relatively hyperbolic groups, of which hyperbolic groups are an instance).

Theorem 6.15: *Let L_1, L_2, \dots be a sequence of Γ -limit groups and π_1, π_2, \dots a sequence of epimorphisms such that $\pi_i : L_i \rightarrow L_{i+1}$. Then all but finitely many of the π_i are isomorphisms.*

Sela also gets as a by-product of the proof of this theorem the following proposition:

Proposition 6.16: *If L is a Γ -limit group obtained as the limit of a sequence of morphisms $h_n : G \rightarrow \Gamma$, then a subsequence of the h_n factors through L .*

We now proceed as in the free limit group case to define shortening quotients.

Definition 6.17: (shortening quotient) *Let L be a Γ -limit group obtained as the limit of a sequence of morphisms $h_n : G \rightarrow \Gamma$. If the h_n are short and **non-injective** we call the limit group obtained a shortening quotient.*

Here again, the shortening argument 4.25 gives

Proposition 6.18: *Let G be a torsion-free freely indecomposable finitely generated non-cyclic group in which every virtually solvable subgroup is free abelian. If L is a shortening quotient of G it is a proper quotient of G .*

We want to show the following result:

Proposition 6.19: *Let G be a torsion-free, freely indecomposable, finitely generated and non-cyclic group. Suppose that every virtually solvable subgroup of G is free abelian. There exists a finite set of Γ -limit groups which are proper quotients of G such that any non-injective morphism f from G to Γ factors through one of the corresponding quotient maps after precomposition by a modular automorphism.*

Remark 6.20: *Note that in this case we need to rule out injective morphisms which obviously do not factor through proper quotients. In the free group case, this was not necessary as no injection can exist from a freely indecomposable non-cyclic group into a free group.*

Here again we introduce the same partial order relation on the shortening quotients of G , and proceed to show that every shortening quotient is under a maximal shortening quotient, and that these are in finite number. The proofs differ slightly from the free case, because we do not have finite presentation of Γ -limit groups. However the result given by proposition 6.16 is enough.

Proposition 6.21: *Let G be a finitely generated group. Every shortening quotient of G is smaller than a maximal shortening quotient.*

Proof. As in the proof of 6.11, we can see that it is enough to show that an infinite countable totally ordered set of shortening quotients has an upper bound. Let thus $(Q_n)_{n \in \mathbb{N}}$ be a set of shortening quotients, with quotient maps η_n . Each Q_n is the limit of a stable sequence $(h_j^n)_{j \in \mathbb{N}}$, we may assume that the kernels of h_j^n and η_n coincide on words of length at most j . Let $\eta : G \rightarrow Q$ be the limit of (a stable subsequence of) the diagonal sequence $(h_n^n)_{n \in \mathbb{N}}$: it is a shortening quotient. We may assume after further extraction that the kernels of η and h_n^n coincide on words of length at most n . Thus the kernels of η and η_n coincide on words of length at most n so that η is the

limit of the sequence $(\eta_n)_{n \in \mathbb{N}}$. By proposition 6.16, for n large enough η_n factors through η . Thus $Q_n < Q$. This terminates the proof. \square

Proposition 6.22: *There is only a finite number of maximal shortening quotient.*

Proof. Assume $(M_n)_{n \in \mathbb{N}}$ is an infinite sequence of distinct maximal shortening quotients of G with quotient map η_n . Each M_n is the limit of a sequence $(h_j^n)_{j \in \mathbb{N}}$ of non-injective short morphisms $G \rightarrow \mathbb{F}$ which we choose to ensure that the kernels of the maps h_j^n and η_n agree on words of length at most j .

As in the proof of 6.21, let $\eta : G \rightarrow M$ be the limit of (a stable subsequence of) the diagonal sequence $(h_j^n)_{n \in \mathbb{N}}$, and see that we may assume that for n large enough, η_n factors through η . Since the M_n are maximal, we have $M = M_n$ for all n large enough, which contradicts the assumption that the maximal quotients M_n are pairwise distinct. \square

We can now prove proposition 6.19:

Proof. Let M_1, \dots, M_r be the maximal proper shortening quotients of G with corresponding quotient maps η_1, \dots, η_r (there are finitely many of them by 6.22). Let f be a **non-injective** morphism $G \rightarrow \mathbb{F}$. Let $\sigma \in \text{Mod}(G)$ be such that $\bar{f} = f \circ \sigma$ is short. The sequence \bar{f}, \bar{f}, \dots gives in the limit a proper shortening quotient Q of G , with quotient map η . Note that \bar{f} factors through η . By proposition 6.21, η factors through one of the η_j . But this implies \bar{f} factors through the same η_j , thus proving the claim. \square

6.3 Relative results

Finally we will show a relative version of proposition 6.19, as well as a 'partial relative version'. We proceed in a similar manner, here the only difference is that we will use the relative version of the shortening argument 4.33.

Let Γ be a hyperbolic torsion-free group with generating set Σ .

Definition 6.23: (relative Γ -limit group) *Let G be a finitely generated group, let H be a subgroup of G , with a fixed embedding into Γ . A Γ -limit group relative to H is the limit of a stable sequence $(h_n)_{n \in \mathbb{N}}$ of homomorphisms $G \rightarrow \Gamma$ which fixes H in the limit (recall definition 4.29). If moreover the homomorphisms h_n are pairwise non-conjugate, we say that L is a strict Γ -limit group relative to H .*

Remark 6.24: *As in the non-relative case, we can see that a Γ -limit group relative to H is either a strict Γ -limit group relative to H , or a subgroup of Γ which contains H .*

A Γ -limit group relative to H is in particular a Γ -limit group, thus we know that a decreasing sequence of Γ -limit groups relative to H stabilises.

Definition 6.25: (relative shortening quotient) *Let L be a Γ -limit group relative to H obtained as the limit of a sequence of morphisms $h_n : G \rightarrow \Gamma$. If the h_n are short relative to H (recall definition 4.32) and non-injective we say that L is a shortening quotient relative to H .*

The relative shortening argument 4.33 gives

Proposition 6.26: *Let G be a finitely generated torsion-free group in which every virtually solvable subgroup is free abelian. Let H be a non-abelian subgroup of G with respect to which G is freely indecomposable. Fix an embedding $H \rightarrow \Gamma$. A shortening quotient of G relative to H is a proper quotient of G .*

Again, we introduce a partial order relation on the shortening quotients of G , and proceed to show that every shortening quotient is under a maximal shortening quotient, and that these are in finite number. The proofs are exactly the same as in the non-relative case, up to checking that the sequences we get do fix H in the limit. We get

Proposition 6.27: *Every shortening quotient of G relative to H is smaller than a maximal shortening quotient relative to H .*

Proposition 6.28: *There is only a finite number of maximal shortening quotient of G relative to H .*

As in the previous sections, this enables us to prove

Proposition 6.29: *Let G be a finitely generated group in which every virtually solvable subgroup is free abelian. Let H be a non-abelian subgroup of G with respect to which G is freely indecomposable. Fix an embedding $H \rightarrow \Gamma$. There exists a finite set of proper Γ -limit quotients of G relative to H such that any non-injective morphism h from G to Γ which fixes H factors through one of the corresponding quotient maps after precomposition by an element of $\text{Mod}_H(G)$.*

But in fact we can prove a slightly more general result

Proposition 6.30: *Let G be a finitely generated group in which every virtually solvable subgroup is free abelian. Let H be a non-abelian subgroup of G with respect to which G is freely indecomposable. Fix an embedding $H \rightarrow \Gamma$. There exists a finite set of proper Γ -limit quotients of G relative to H , and a finite subset H_0 of H such that any non-injective morphism h from G to Γ which fixes H_0 factors through one of the corresponding quotient maps after precomposition by an element of $\text{Mod}_H(G)$.*

Proof. Let L_1, \dots, L_p be the maximal shortening quotients of G relative to H with corresponding quotient maps η_1, \dots, η_p (there are finitely many of them by 6.28). Suppose there exists no such H_0 . Then we can produce a sequence $h_n : G \rightarrow \Gamma$ of non-injective morphisms which are short relative to H , fix H in the limit, and such that none of the maps h_n factors through any of the η_j . From this sequence extract a stable sequence, which converges to a shortening quotient Q of G relative to H .

This quotient Q is under one of the maximal shortening quotients L_j so that the quotient map $\pi : G \rightarrow Q$ factors through η_j . Now by proposition 6.16, an infinity of the h_n factor through π , and thus through η_j . This is a contradiction, and we have completed the proof. \square

Now recall that Theorem 4.38 tells us that if G is a torsion-free hyperbolic group freely indecomposable with respect to a subgroup H , an injective morphism $G \rightarrow G$ which fixes a large enough subset of H has to be surjective. Thus we get as an immediate corollary

Corollary 6.31: *Let G be a torsion-free hyperbolic group, and let H be a non-abelian subgroup of G with respect to which G is freely indecomposable. There exist a finite set of proper quotients of G , and a finite subset H_0 of H such that any non-surjective morphism $h : G \rightarrow G$ which fixes H_0 factors through one of these quotients after precomposition by an element of $\text{Mod}_H(G)$.*

Chapter 7

Elementary embeddings in a hyperbolic group

7.1 Hyperbolic towers and statement of the main result

The surfaces we consider are always compact and connected. We define hyperbolic towers.

Definition 7.1: (hyperbolic floor) *Consider a triple (G, G', r) where G is a group, G' is a subgroup of G , and r is a retraction from G onto G' . We say that (G, G', r) is a hyperbolic floor if there exists a non-trivial decomposition Λ of G as a graph of groups with surfaces (recall definition 4.12) such that:*

- *the graph of groups Λ has exactly one vertex w which is not of surface type, and its vertex group is G' ;*
- *every edge of Λ is adjacent to w ;*
- *the endpoints of an edge are distinct;*
- *for each vertex v distinct from w , the image of G_v by the retraction r is non-abelian.*

Definition 7.2: (hyperbolic tower) *Let G be a group, let H be a subgroup of G . We say that G is a hyperbolic tower based on H if there exists a finite sequence $G = G^0 > G^1 > \dots > G^m > H$ of subgroups of G where:*

- *for each k in $[0, m-1]$, there exists a retraction $r_k : G^k \rightarrow G^{k+1}$ such that the triple (G^k, G^{k+1}, r_k) is a hyperbolic floor.*
- *$G^m = H * F * S_1 * \dots * S_p$ where F is a (possibly trivial) free group, $p \geq 0$, and each S_i is the fundamental group of a closed and connected surface of Euler characteristic at most -2 ;*

Remark 7.3: *If G is a hyperbolic tower over a subgroup H , and G' is a hyperbolic tower over a subgroup H' , then $G * G'$ is a hyperbolic tower over $H * H'$. If G is a hyperbolic tower over a subgroup G' , and G' is a hyperbolic tower over a subgroup H , then G is a hyperbolic tower over H .*

We can now state our main theorem.

Theorem 7.4: *Let G be a torsion-free hyperbolic group. Let $H \hookrightarrow G$ be an elementary embedding. Then G is a hyperbolic tower based on H .*

This implies in particular that H is finitely generated, and a retract of G .

7.2 JSJ-like decompositions and preretractions

To prove Theorem 7.4, we need to construct successive retractions from subgroups of G to proper subgroups until we get to H . The strategy will be to build, by the mean of first-order sentences, some maps that we will call preretractions: their properties will allow us to build the retractions we need. These preretractions are associated to a specific type of graphs of groups.

7.2.1 JSJ-like decomposition

Definition 7.5: (JSJ-like decomposition) *Let Λ be a graph of groups with surfaces whose edge groups are infinite cyclic. Let A be the fundamental group of Λ . Call Z type vertices the vertices of Λ which are not of surface type and have infinite cyclic vertex group, and rigid type vertices the other non surface type vertices. We say that Λ is a cyclic JSJ-like decomposition of A if:*

1. *an edge is adjacent to at most one surface type vertex, and to at most one Z type vertex;*
2. *(strong 2-acylindricity) if a non-trivial element of A stabilises two distinct edges of T_Λ , they are adjacent and their common endpoint is the lift of a Z type vertex.*

Remark 7.6: *These conditions imply in particular that two distinct vertices have distinct vertex groups.*

Remark 7.7: *Let Λ be a cyclic JSJ-like decomposition of a group A . If C is the edge group of an edge e which connects two vertices which are not of Z type, then C is maximal abelian in A . Indeed if c is a non-trivial element of C , and if γ commutes with c , then c fixes the edge $\gamma \cdot e$. By strong 2-acylindricity, c does not fix any edges other than e , so $\gamma \cdot e = e$, and thus $\gamma \in C$.*

We will also say that a vertex in the tree T_Λ is of type Z or rigid according to the type of its image by the quotient map $T_\Lambda \rightarrow \Lambda$.

Remark 7.8: *Note that a rigid type vertex group in a JSJ-like decomposition might admit some splittings over \mathbb{Z} compatible with Λ , so is not rigid in the usual sense of the term.*

Definition 7.9: (subgroups with disjoint conjugacy classes) *We say that two subgroups of A have disjoint conjugacy classes if no non-trivial element of one of the subgroups has a conjugate in the other.*

Remark 7.10: *Given a strongly 2-acylindrical graph of groups decomposition Λ of a group A , consider two edge stabilisers G_1 and G_2 of the tree T_Λ . Denote by e_1 and e_2 the projection in Λ of the edges they stabilise. We claim that G_1 and G_2 have disjoint conjugacy classes unless either $e_1 = e_2$, or e_1 and e_2 are adjacent to a common Z type vertex. This is an easy consequence of strong 2-acylindricity.*

7.2.2 JSJ decompositions

A JSJ decomposition Λ of a group G is a decomposition as a graph of groups which encodes all possible splittings of the group G over a given class \mathcal{E} of subgroups. The standard reference for the case where G is finitely generated and one ended are Rips and Sela in [RS97], see also [DS99] and [FP06].

In the sequel, we will use the JSJ decomposition in the case where G is torsion-free hyperbolic and freely indecomposable (respectively freely indecomposable with respect to a subgroup H), and \mathcal{E} is the class of cyclic groups. We call such a decomposition a cyclic JSJ decomposition of G (respectively a relative cyclic JSJ decomposition with respect to H).

In both these casee, a cyclic JSJ decomposition Λ of G is a graph of groups with surfaces. Moreover, we can see from Theorem 7.1 in [RS97] that for any other decomposition Γ of G as a graph of groups with surfaces and cyclic edge groups we have:

- (i) non surface type vertex groups of Λ are elliptic in Γ ;
- (ii) edge groups of Λ are elliptic in Γ ;
- (iii) surface type vertex groups of Γ are contained in a surface type vertex group of Λ .

We will need two properties of such a cyclic (relative) JSJ decomposition Λ : the first, explained in the following remark, is that we may assume that it is a JSJ-like decomposition. The second is that its vertex groups are 'preserved' under modular automorphisms, namely that an element of $\text{Mod}(G)$ (respectively $\text{Mod}_H(G)$) restricts to a conjugation on non surface type vertex groups of Λ , and sends surface type vertex groups of Λ isomorphically on conjugates of themselves: this will be the object of lemma 7.12.

Remark 7.11:

- *Let G be a torsion-free hyperbolic group. Suppose G is freely indecomposable. Let Λ be the JSJ decomposition of G given by Theorem 7.1 of [RS97]. Then we may assume Λ is JSJ-like.*

Let g be a non-trivial primitive element of G , and denote T_g the subtree whose edges have stabilisers lying in $\langle g \rangle$. We claim that the translates of T_g are all disjoint. Indeed, suppose that there is an edge e which lies both in T_g and in $h \cdot T_g$ for some element h of G . This implies that some power g^j of g fixes e , and that $e = h \cdot e'$ for some edge e' in T_g . Since e' is fixed by some power g^k of g , we get that e is also fixed by hg^kh^{-1} . Edge stabilisers are abelian, so $[g^j, hg^kh^{-1}] = 1$. Now in torsion-free hyperbolic groups, maximal abelian subgroups are malnormal, so h is also a power of g , and $T_g = h \cdot T_g$. We now remove in T the interior of all the edges of T_g , add a vertex v_g and edges from each vertex of T_g to the new vertex v_g . Since the translates of T_g are disjoint, we can do this equivariantly. It is then fairly straightforward to see that the modified decomposition still satisfies the properties (i), (ii) and (iii) above, and is strongly 2-acylindrical.

- *Let G be a torsion-free hyperbolic group which is freely indecomposable with respect to a subgroup H . Let Λ be the cyclic relative JSJ decomposition of G with respect to H given by Théorème 4.1 of [Pau03]. Similarly, we can assume that Λ is a cyclic JSJ-like decomposition.*

The following lemma describes the other property of cyclic JSJ and relative cyclic JSJ of a hyperbolic group that we will need. It is a consequence of the universal properties of the JSJ decomposition, i.e. of the fact that it describes any cyclic splitting of the group.

Lemma 7.12: *Let G be a torsion-free hyperbolic group which is freely indecomposable. Denote by Λ its cyclic JSJ decomposition, as given by Theorem 7.1 of [RS97]. An element of $\text{Mod}(G)$ restricts to conjugation on each non surface type vertex group of Λ , and sends surface type vertex groups isomorphically on conjugates of themselves.*

Similarly, suppose G is a torsion-free hyperbolic group which is freely indecomposable with respect to a subgroup H . Let Λ denote its cyclic relative JSJ decomposition with respect to H . An element of $\text{Mod}_H(G)$ restricts to conjugation on each non surface type vertex group of Λ , and sends surface type vertex groups isomorphically on conjugates of themselves.

Proof. Let Γ be a decomposition of G as a graph of groups with surfaces with cyclic edge stabilisers.

Let S be a surface type vertex group of the cyclic JSJ decomposition Λ . The boundary elements of S stabilise edges of Λ , thus by property (ii) above, they stabilise vertices of Γ . Since S acts faithfully on Γ in such a way that its boundary elements are elliptic, and with cyclic edge stabilisers, by lemma III.2.6 of [MS84] it inherits a decomposition Γ_S which is dual to a set of non boundary parallel simple closed curves on the corresponding surface, such that there is an

equivariant injective simplicial map $t : T_{\Gamma_S} \rightarrow T_\Gamma$. In particular, every vertex group of Γ_S is elliptic in Γ .

Let S_0 be a vertex group of Γ_S . If S_0 lies in a surface type vertex group S' of Γ , we claim that $S_0 = S'$. Indeed, by property (iii) above, S' lies in a surface type vertex group of Λ . If this group is not S , this means S_0 stabilise an edge of Λ , so that it is a boundary subgroup of S . But this contradicts the fact that the simple closed curves are not boundary parallel. Thus S' is a subgroup of S . On the other hand, the injectivity of the map t shows that $S \cap S' = S_0$. This proves the claim. If S_0 lies in a non surface type vertex group of Γ , by definition of the modular group, a modular automorphism restricts to conjugation on S_0 .

Now we see that in both cases, a modular automorphism σ of Γ sends S_0 isomorphically on a conjugate of itself. Since the graph of groups Γ_S is connected, and its edge groups are non-trivial, this implies that σ sends S isomorphically on a conjugate of itself.

Consider now the case of a non surface type vertex group R of Λ : by property (i) above, it is elliptic in Γ . If R lies in a non surface type vertex group R' of Γ , the restriction of any element of $\text{Mod}(\Gamma)$ to R' , and thus to R , is just a conjugation. Suppose on the other hand that R lies in a surface type vertex group S_R of Γ . By part (iii) of Theorem 7.1 of [RS97], S_R itself lies in a surface type vertex group S of Λ . But by our argument above, boundary subgroups of S are sent on conjugates by modular automorphisms. \square

7.2.3 Preretractions

Preretractions are maps that preserve some of the structure of a JSJ-like decomposition. We need to define them as maps $A \rightarrow G$ where A is a subgroup of G .

Definition 7.13: (preretraction) *Let G be a torsion-free hyperbolic group. Let A be a subgroup of G , and Λ a JSJ-like decomposition of A . A map $A \rightarrow G$ is a preretraction with respect to Λ if its restriction to each non surface type vertex group A_v of Λ is just conjugation by some element g_v of G , and if surface type vertex groups have non-abelian images.*

Remark 7.14: *The definition of a JSJ-like decomposition implies that the restriction of a preretraction to an edge group is just conjugation by an element of G . Indeed by condition 1 of definition 7.5, every edge group is contained in at least one non surface type vertex group.*

In the next two sections, we will prove Theorem 7.4, using two results about preretractions. The last two chapters of this thesis are devoted to their proofs: they are both intermediate steps in the proof of proposition 6 of [Sel06] but are not explicitly stated there. The first is given by

Proposition 7.15: *Let A be a torsion-free hyperbolic group. Let Λ be a cyclic JSJ-like decomposition of A . Assume that there exists a non-injective preretraction $A \rightarrow A$ with respect to Λ . Then there exists a subgroup A' of A , and a retraction r from A to A' , such that (A, A', r) is a hyperbolic floor. Moreover, given a rigid type vertex group R_0 of Λ , we can choose r such that R_0 is in the image of r .*

The second proposition is needed to complete the induction step in the construction of our hyperbolic tower.

Proposition 7.16: *Let G be a torsion-free hyperbolic group. Let A be a group which admits a JSJ-like decomposition Λ . Suppose G' is a subgroup of G containing A such that either G' is a free factor of G , or G' is a retract of G by a retraction $r : G \rightarrow G'$ which makes (G, G', r) a hyperbolic floor. If there exists a non-injective preretraction $A \rightarrow G$ with respect to Λ , then there exists a non-injective preretraction $A \rightarrow G'$ with respect to Λ .*

7.3 Using first order to build preretractions

Suppose H is a subgroup elementarily embedded in a torsion-free hyperbolic group G . To show that G admits a structure of hyperbolic tower over H , we will start by decomposing G in free factors relatively to H . That is, we will write $G = A * B_1 \dots * B_m$ where the groups B_j are freely indecomposable (possibly infinite cyclic) and A is freely indecomposable with respect to H . We call such a decomposition a Grushko decomposition of G relative to H .

If we can show that A admits a structure of hyperbolic tower over H , and that the groups B_i admit a structure of hyperbolic tower over 1, we will be done by remark 7.3. The idea is thus to produce non-injective preretractions $A \rightarrow A$ and $B_i \rightarrow B_i$, in order to be able to apply the propositions 7.15 and get the top floor of a hyperbolic tower decomposition. But for this, it is enough by proposition 7.16 to build non-injective preretractions $A \rightarrow G$ and $B_i \rightarrow G$. This is what the following two propositions will enable us to do. In fact they are slightly more general. This greater generality is required for the induction step, when we will build further floors of our hyperbolic towers.

The heart of the proof of Theorem 7.4 is contained in proposition 7.20 and proposition 7.21. The idea is that by expressing the existence of a factor set in first-order logic, we can prove the existence of a non-injective preretraction.

We will use the following definition

Definition 7.17: (Λ -related morphisms) *Let A be a group which admits a JSJ-like decomposition Λ . We say two morphisms f and f' from A to a group G are Λ -related if*

- *for each non surface type vertex group R of Λ , there exists an element u_R such that f' restricted to R is $\text{Conj}(u_R) \circ f$;*
- *each surface type vertex group of Λ that has non-abelian image by f also has non-abelian image by f' .*

Remark 7.18: *Note that if A is a subgroup of G , a map $f : A \rightarrow G$ is Λ -related to the embedding $A \hookrightarrow G$ if and only if it is a preretraction.*

The following lemma shows that relatedness can be expressed in first-order logic.

Lemma 7.19: *Let A be a group finitely generated by a tuple \mathbf{a} . Suppose A admits a JSJ-like decomposition Λ . There exists a formula $\text{Rel}(\mathbf{x}, \mathbf{y})$ such that for any pair of morphisms f and f' from A to G , the formula $\text{Rel}(f(\mathbf{a}), f'(\mathbf{a}))$ is satisfied by G if and only if f and f' are Λ -related.*

Proof. We introduce some notation. Denote by R_1, \dots, R_r the non surface type vertex groups of Λ , and by S_1, \dots, S_s its surface type vertex groups. For $1 \leq p \leq r$, choose a finite generating set ρ_p for R_p , and for $1 \leq q \leq s$, choose a finite generating set σ_q for S_q . We take the convention to denote tuples by bold font, and to denote by $l(\mathbf{x})$ the cardinality of the tuple \mathbf{x} .

The elements of σ_p and ρ_q can be represented by words in the elements \mathbf{a} , we denote these by $\sigma_p = \bar{\sigma}_p(\mathbf{a})$ and $\rho_q = \bar{\rho}_q(\mathbf{a})$ respectively.

Two maps f and f' satisfy the condition on the rigid type vertex groups of Λ if and only if

$$\exists u_1 \dots u_r \bigwedge_{p=1}^r \{f'(\rho_p) = u_p f(\rho_p) u_p^{-1}\}.$$

To express the abelianity of a subgroup generated by a tuple $\mathbf{z} = (z^1, \dots, z^{l(\mathbf{z})})$, we can use the formula $\text{Ab}(\mathbf{z}) : \bigwedge_{i,j} \{[z^i, z^j] = 1\}$. Thus the non-abelianity condition about the image by f and f' of surface type vertex groups of Λ can also be expressed by

$$\bigwedge_{q=1}^s \{\neg \text{Ab}(f(\sigma_q)) \rightarrow \neg \text{Ab}(f'(\sigma_q))\}.$$

But now, if w is an element of A which can be represented by a word $\bar{w}(\mathbf{a})$, its image by the morphism extending $\mathbf{a} \mapsto \mathbf{x}$ is represented by $\bar{w}(\mathbf{x})$.

Thus if $f : \mathbf{a} \rightarrow \mathbf{x}$ and $f' : \mathbf{a} \rightarrow \mathbf{y}$, the formula $\text{Rel}(\mathbf{x}, \mathbf{y})$ with free variables \mathbf{x}, \mathbf{y} given by

$$\exists u_1 \dots u_r \left[\bigwedge_{p=1}^r \{ \bar{\rho}_p(\mathbf{y}) = u_p \bar{\rho}_p(\mathbf{x}) u_p^{-1} \} \right] \wedge \left[\bigwedge_{q=1}^s \{ \neg \text{Ab}(\bar{\sigma}_q(\mathbf{x})) \rightarrow \neg \text{Ab}(\bar{\sigma}_q(\mathbf{y})) \} \right]$$

is satisfied by G if and only if f and f' are Λ -related. □

We can now prove the two key propositions.

Proposition 7.20: *Suppose that G is a non-cyclic torsion-free hyperbolic group, and let H be a subgroup elementarily embedded in G . Suppose A is a hyperbolic subgroup of G which properly contains H , and which is freely indecomposable relative to H . Let Λ be the cyclic JSJ decomposition of A relative to H . Then there exists a non-injective preretraction $A \rightarrow G$ with respect to Λ .*

Proof. Corollary 6.31 applied to A tells us that there exist a finite subset H_0 of H , and a finite family of proper quotients $\eta_j : A \rightarrow L_j$ for $j \in [1, m]$, such that any non-surjective morphism $\theta : A \rightarrow A$ which fixes H_0 factors through one of the quotients η_j after precomposition by an element of $\text{Mod}_H(A)$.

A morphism $\theta : A \rightarrow H$ can be seen as a non-surjective morphism $A \rightarrow A$ since we assumed $H \neq A$. Thus any morphism $\theta : A \rightarrow H$ which fixes H_0 factors through one of the quotients η_j after precomposition by an element τ of $\text{Mod}_H(A)$.

We will give a weaker form of this statement, in order to be able to express it as a first order sentence satisfied by H . Indeed, since we cannot express the modular group with first order logic, we have to lose some information.

For each l in $[1, m]$, we fix an element ν_j in the kernel of $\eta_j : A \rightarrow Q_l$. Let Λ be the cyclic JSJ decomposition of A relative to H . By lemma 7.12, elements of $\text{Mod}_H(A)$ restrict to a conjugation on non-surface type vertex groups of Λ , and send surface type vertex groups isomorphically on conjugates of themselves.

Thus, if θ is a morphism $A \rightarrow H$, and if $\tau \in \text{Mod}_H(A)$, the restriction of the map $\theta' = \theta \circ \tau$ to each rigid type vertex group coincides with θ up to conjugation, and if the image by θ of a surface type vertex group S of Λ is non-abelian, so is the image of S by θ' . This says exactly that θ and θ' are Λ -related.

This implies that the following statement holds:

Statement 1: *For any morphism $\theta : A \rightarrow H$ which fixes H_0 , there exists a morphism $\theta' : A \rightarrow H$ such that θ and θ' are Λ -related, and there exists j in $[1, m]$ such that $\theta'(\nu_j) = 1$.*

Let us now see that this statement can be expressed by a first order sentence in the language \mathcal{L}_H which is satisfied by H .

The group A is hyperbolic, we choose a finite presentation $\langle \mathbf{a} \mid \bar{\Sigma}_A(\mathbf{a}) \rangle$. If an $l(\mathbf{a})$ -tuple \mathbf{x} in H satisfies $\bar{\Sigma}_A(\mathbf{x}) = 1$, the map $A \rightarrow H$ which sends \mathbf{a} to \mathbf{x} is a morphism. Conversely, any morphism $A \rightarrow H$ comes from a solution to the system of equations $\bar{\Sigma}_A(\mathbf{x}) = 1$ in H . The elements ν_l can be represented by words $\bar{\nu}_l(\mathbf{a})$; and for each h in H_0 , the element h can be represented by a word $\bar{h}(\mathbf{a})$.

Recall that the language \mathcal{L}_H is defined as the language of groups to which we have added a constant symbol $[h]$ for each h in H . To express that the morphism corresponding to the tuple \mathbf{x} fixes the finite subset H_0 of H , we can thus write

$$\bigwedge_{h \in H_0} \{ [h] = \bar{h}(\mathbf{x}) \}.$$

To express that the morphism corresponding to the tuple \mathbf{x} sends one of the elements ν_i to 1, we can write

$$\bigvee_{i=1}^m \{\bar{\nu}_i(\mathbf{x}) = 1\}.$$

Finally consider the sentence (\dagger) over \mathcal{L}_H given by

$$\forall \mathbf{x} \left[\Sigma_A(\mathbf{x}) = 1 \wedge \bigwedge_{h \in H_0} [h] = \bar{h}(\mathbf{x}) \right] \rightarrow \exists \mathbf{y} \left[\bar{\Sigma}_A(\mathbf{y}) = 1 \right] \wedge \text{Rel}(\mathbf{x}, \mathbf{y}) \wedge \left[\bigvee_{l=1}^m \bar{\nu}_l(\mathbf{y}) = 1 \right].$$

We claim that the interpretation of the first-order formula (\dagger) on H is exactly Statement 1, so it is true on H . To see this, let us interpret (\dagger) on H . The formula in the square brackets on the first line of (\dagger) says that the $l(\mathbf{x})$ -tuple \mathbf{x} satisfies $\bar{\Sigma}(\mathbf{x}) = 1$, so that the map $\mathbf{x} \mapsto \mathbf{a}$ extends to a morphism $\theta : A \rightarrow H$. Moreover, for each h in H_0 , we have $h = \bar{h}(\mathbf{x}) = \theta(h)$, so θ fixes H_0 (recall that the interpretation of the constant $[h]$ on H is just the element h). The next part of the formula says that there exists a $l(\mathbf{y})$ -tuple \mathbf{y} satisfying $\bar{\Sigma}(\mathbf{y}) = 1$, and whose corresponding morphism $\theta' : A \rightarrow H$ is related to θ . Finally the formula in the last square brackets says that for at least one value of l , we have $\bar{\nu}_l(\mathbf{y}) = \theta'(\nu_l) = 1$. That is, at least one of the elements ν_l is in the kernel of θ' . This proves the claim.

The formula (\dagger) is therefore satisfied by G . Recall that it can be interpreted on G : the symbols we added to the language of groups are constants $[h]$ for each element h of H , and $H \leq G$ so we just interpret $[h]$ by h . If we take the 'tautological solution' \mathbf{a} to the equation $\Sigma_A(\mathbf{x}) = 1$, it satisfies the formula in the first square brackets: indeed, $\Sigma_A(\mathbf{a}) = 1$, and for each $h \in F$, we have $h = \bar{h}(\mathbf{a})$ by definition of \bar{h} . Thus we get a tuple \mathbf{y} such that $\mathbf{a} \mapsto \mathbf{y}$ extends to a morphism μ , which is Λ -related to the morphism $\mathbf{a} \mapsto \mathbf{a}$. Since it sends one of the elements ν_i to 1, it is not injective. But the morphism $\mathbf{a} \mapsto \mathbf{a}$ is just the embedding $A \hookrightarrow G$, so by remark 7.18, $\mu : A \rightarrow G$ is a non-injective preretraction. \square

We now show the second key result.

Proposition 7.21: *Suppose that G is a torsion-free hyperbolic group, and that H is a subgroup elementarily embedded in G which is also a retract of G . Let B be a freely indecomposable hyperbolic subgroup of G which is neither cyclic nor a closed surface group of Euler characteristic at most -2 . Let Λ be the cyclic JSJ decomposition of B .*

Suppose that no non-trivial element of B is conjugate in G to an element of H . Then there exists a non-injective preretraction $B \rightarrow G$ with respect to Λ .

Proof. Assume first that B is not the fundamental group of the closed surface of Euler characteristic -1 . We choose a presentation $\langle \mathbf{b} \mid \bar{\Sigma}_B(\mathbf{b}) \rangle$ for B . Let Λ be the cyclic JSJ decomposition of B .

Let $\eta_1 : B \rightarrow L_1, \dots, \eta_m : B \rightarrow L_m$ be the proper quotients of B given by proposition 6.19. Again we choose elements ν_1, \dots, ν_m of B such that ν_j is in the kernel of η_j .

Proposition 6.19 tells us that any non-injective map from B to G factors through one of the quotients η_j after precomposition by an element of $\text{Mod}(B)$. Note that a map $\theta : B \rightarrow H$ can be seen as a map $B \rightarrow G$, so the previous statement remains true if we replace $B \rightarrow G$ by $B \rightarrow H$.

We want to find a sufficient condition for non-injectivity of a map $B \rightarrow H$ that is expressible in first-order. Proposition 4.34, applied to B with $H = 1$, tells us that there exist a finite set i_1, \dots, i_t of embeddings of B in G such that for any embedding $i : B \hookrightarrow G$, there exists an element σ of $\text{Mod}(B)$, an integer k in $[1, t]$ and an element g of G such that

$$i(x) = gi_k(\sigma(x))g^{-1} \text{ for all } x \in B$$

Remark that if i is an embedding of B in H , we can apply the retraction $r : G \rightarrow H$ given by the hypotheses to both sides, and note that σ is an isomorphism, to get that there exists an element τ of $\text{Mod}(B)$, an integer k in $[1, t]$ and an element g of G such that for all x in B ,

$$r(g^{-1})i(\tau(x))r(g) = r(i_k(x))$$

By lemma 7.12, the map on the left hand side is Λ -related to i . So if i is an embedding $B \hookrightarrow H$, there exists a map i' which is Λ -related to i , and which satisfies $i'(\mathbf{b}) = r(i_k(\mathbf{b}))$ for some k .

Let $\theta : B \rightarrow H$. Consider the following statement about θ , that we denote $S(\theta)$.

$S(\theta)$: Suppose that $\theta' : B \rightarrow H$ is a morphism which is Λ -related to θ . Then for any integer k in $[1, t]$, we have $\theta'(\mathbf{b}) \neq r(i_k(\mathbf{b}))$.

From the previous paragraph, if $S(\theta)$ holds, then θ isn't an embedding: it is a sufficient condition for a map not to be an embedding.

Again by lemma 7.12, if θ is a morphism $B \rightarrow H$, and if $\tau \in \text{Mod}(B)$, the maps $\theta' = \theta \circ \tau$ and θ are Λ -related. So the following statement is true.

Statement 2: If $\theta : B \rightarrow H$ is a morphism such that $S(\theta)$ holds, then there exists a morphism $\theta'' : B \rightarrow H$ and l in $[1, m]$ such that θ and θ'' are Λ -related, and $\theta''(\nu_l) = 1$.

This is the statement we want to express by a first-order formula. Let us first try to see that $S(\theta)$ can be expressed by a first order formula on the variables $\theta(\mathbf{b})$. Consider the following first order formula $\psi(\mathbf{x})$ with free variable the $l(\mathbf{b})$ -tuple \mathbf{x} :

$$[\bar{\Sigma}_B(\mathbf{x}) = 1] \wedge \forall \mathbf{z} [\bar{\Sigma}_B(\mathbf{z}) = 1 \wedge \text{Rel}(\mathbf{z}, \mathbf{x})] \rightarrow \left[\bigwedge_{k=1}^t \mathbf{z} \neq [r(i_k(\mathbf{b}))] \right].$$

This is a first order formula in the language \mathcal{L}_H . Thus the constant $[r(i_k(\mathbf{b}))]$ is interpreted in both H and G simply by the element $r(i_k(\mathbf{b}))$ (which is indeed an element of H).

Let \mathbf{x} be a $l(\mathbf{b})$ -tuple in H . It is straightforward to see that the formula $\psi(\mathbf{x})$ is satisfied by H if and only if the map $\theta : \mathbf{b} \mapsto \mathbf{x}$ is a morphism for which the statement $S(\theta)$ holds. So if $\psi(\mathbf{x})$ is satisfied by H , the map $\theta : \mathbf{b} \mapsto \mathbf{x}$ is a non-injective morphism $B \rightarrow H$.

We can now write the first order sentence $(\dagger\dagger)$

$$\forall \mathbf{x} \psi(\mathbf{x}) \rightarrow \exists \mathbf{y} [\bar{\Sigma}_B(\mathbf{y}) = 1] \wedge \text{Rel}(\mathbf{x}, \mathbf{y}) \wedge \left[\bigvee_{j=1}^l \bar{\nu}_j(\mathbf{y}) = 1 \right].$$

Just as we saw that (\dagger) expressed Statement 1 in 7.20, we can see that the first order formula $(\dagger\dagger)$ on H expresses Statement 2, so it is satisfied by H .

As H is elementarily embedded in G , the formula $(\dagger\dagger)$ is also satisfied by G . As in the proof of 7.20, we can apply it to the tautological solution \mathbf{b} of $\bar{\Sigma}_B(\mathbf{x}) = 1$. To see that $G \models \psi(\mathbf{b})$, note first that by our hypotheses, the JSJ decomposition of B admits at least one non-surface type vertex group. A map $\mu : B \rightarrow G$ which is Λ -related to the embedding $\mathbf{b} \mapsto \mathbf{b}$ restricts to conjugation on the non-surface type vertex groups of B , thus it cannot have image in H since no element of B can be conjugated into H by an element of G . This implies in particular that for all k , the $l(\mathbf{b})$ -tuple $\mu(\mathbf{b})$ is distinct from the tuple $r(i_k(\mathbf{b}))$.

The second part of the sentence $(\dagger\dagger)$ thus gives a morphism $B \rightarrow G$ which is Λ -related to the embedding $B \hookrightarrow G$ and kills one of the elements ν_i : it is a non-injective preretraction.

In the case where B is the fundamental group of the non-orientable surface of Euler characteristic -1 , we can follow the same proof if we consider the JSJ of B to consist of a single rigid vertex, and that $\text{Mod}(B)$ is trivial. Indeed, the group of automorphisms of B is finite, so if we

replace the finite list $(\eta_i)_i$ by the finite list $(\eta_i \circ \tau)_{i, \tau \in \text{Aut}(B)}$, the conclusion of 6.19 still hold with our new definition of the modular group. Similarly, the existence of a finite set of embeddings up to conjugation holds. \square

7.4 Proof of the main result

We can now prove Theorem 7.4.

Proof. Let us first treat the case where G is infinite cyclic, generated by an element g . Any subgroup of G is of the form $H = \langle g^m \rangle$. Now G satisfies the formula $\exists x \{x^m = [g^m]\}$, which expresses that g admits an m -th root in G . This is a formula over \mathcal{L}_H which is true on H if and only if g has an m -th root in H , that is if and only if $H = G$. Thus if G is cyclic, its only elementarily embedded subgroup is itself, and the theorem is trivial.

So assume that G is a non-cyclic torsion-free hyperbolic group, and let H be a subgroup elementarily embedded in G . Note that H is necessarily non-abelian as it is elementary equivalent to G .

We will first show that G admits a structure of hyperbolic tower over a group G' whose Grushko decomposition relative to H is of the form $G' = H * B'_1 * \dots * B'_r$. Set $G^0 = G$. We will define by induction a finite sequence $G = G^0 > G^1 > \dots > G^N$ of subgroups of G , such that H is a free factor of G^N , and G^m has a structure of hyperbolic floor over G^{m+1} for each m up to N .

Assume G^m is defined, and write the Grushko decomposition of G_m relative to H as

$$G^m = A^m * B_1^m * \dots * B_{p_m}^m$$

where A^m is the factor containing H . If $A^m = H$ we are done, so assume $A^m \neq H$.

Note that A^m is freely indecomposable relative to H . Denote by Λ the cyclic JSJ of A^m relative to H . Note also that A^m is a retract of G , so it is a quasiconvex subgroup of G , and thus it is hyperbolic.

All the hypotheses of proposition 7.20 for $A = A^m$ are satisfied, so we can apply it to get a non-injective preretraction $A^m \rightarrow G$ with respect to Λ . We now apply proposition 7.16 successively to the floors of the hyperbolic tower formed by G over A^m to get a non-injective preretraction $A^m \rightarrow A^m$ with respect to Λ . Finally by proposition 7.15, we get a retraction $r : A^m \rightarrow A_0^m$ on a proper subgroup of A^m such that (A^m, A_0^m, r) is a floor of a hyperbolic tower and the rigid group of Λ which contains H is in A_0^m . Now define G^{m+1} by

$$G^{m+1} = A_0^m * B_1^m * \dots * B_{p_m}^m.$$

Since A^m has a structure of hyperbolic floor over A_0^m , by remark 7.3, the group G^m has a structure of hyperbolic floor over G^{m+1} as required.

As each G^{m+1} is a strict retract of G^m , and since the groups G^m are all subgroups of G , they are G -limit groups. Thus by proposition 6.15 the sequence is finite. At the end of this process, we get a group G^N in which H is a free factor, and such that G is built as a hyperbolic tower based on G^N .

If all the other factors of the Grushko decomposition of G^N relative to H are surface groups or free groups, we are done. So assume that there is a factor B which is neither free nor a closed surface group. Note that as a retract of G , the group G^N is hyperbolic, so as a free factor of G^N , the group B is itself hyperbolic. We will now show that B has a structure of hyperbolic tower over 1.

Any two conjugates of H and B in G^N intersect trivially, since they are free factors in G^N . But since G^N is a retract of G , any two conjugates of H and B in G must also intersect trivially. Hence the conditions of 7.21 are satisfied by B : by applying it, we get a non-injective preretraction

$B \rightarrow G$. We apply 7.16 iteratively to get a non-injective preretraction $B \rightarrow B$, which by 7.15 gives us a retraction $r : B \rightarrow B'$, such that (B, B', r) is a hyperbolic floor.

Note that since B' is a retract of G , the number of factors in its Grushko decomposition is bounded above by the rank of G . If any of the factors of the Grushko decomposition of B' are neither free nor surface, we can repeat the process above. This terminates, as before, because the groups involved are G -limit groups and because the number of factors in the Grushko decomposition of our groups is bounded. We finally get that B is a hyperbolic tower over 1.

Thus all the factors of G^N distinct from H are hyperbolic towers over 1. By remark 7.3, the group G^N is a hyperbolic tower over H . To conclude, apply once more remark 7.3 to see that G is a hyperbolic tower over H . \square

7.5 The special case of free groups

In the special case where our hyperbolic group is free, Theorem 7.4 together with Theorem 4 in [Sel06] gives

Corollary 7.22: *Let F be a finitely generated free group, let H be a subgroup of F . The embedding of H in F is elementary if and only if H is a free factor of F .*

Proof. Suppose that H is an elementary subgroup of F . By proposition 7.4, F has a structure of hyperbolic tower over H . If the tower has at least one floor, there exists a subgroup F' of F , and a retraction $r : F \rightarrow F'$ so that $H < F'$, and (F, F', r) is a hyperbolic floor built by adding a (possibly disconnected) surface Σ . Let $\gamma_1, \dots, \gamma_r$ be generators of pairwise non-conjugate maximal boundary subgroups of $S = \pi_1(\Sigma)$. We know, from the standard presentation of a surface group with boundary, that the product of the elements γ_i is equal to a product of commutators and squares. Both F and F' being free groups, lemma 4.1 in [BF] tells us that there is a decomposition of F' as $Z * F''$, where Z is an infinite cyclic group, such that one of these boundary subgroups generators, say γ_1 , is a generator of Z , and all the other boundary subgroups generators γ_i are in conjugates of F'' . Now let $\alpha : F' \rightarrow Z/2Z$ be the map which kills F'' and the squares in Z . The image of γ_1 by $\alpha \circ r$ is the generator of $Z/2Z$, and for $i \neq 1$, the image of γ_i is trivial. However, the image of squares and commutators are sent to 1 by $\alpha \circ r$, this is a contradiction. This shows that the only structure of hyperbolic tower that a free group can have over one of its subgroup is a trivial one, where the subgroup is a free factor of the free group. Thus H is a free factor of F . Conversely, if H is a free factor of F , its embedding into F is elementary by Theorem 4 of [Sel06]. \square

Chapter 8

A property of JSJ-like decompositions

To complete the proof of 7.4, we now need to prove proposition 7.15 and 7.16. This will be done in the last chapter, using the results that we will expose in this chapter and the next.

This section aims to show that if a preretraction $G \rightarrow G$ relative to some cyclic JSJ-like decomposition of G satisfies some strong injectivity conditions on the vertex groups, it must be an isomorphism. Recall that a preretraction $A \rightarrow A$ with respect to a JSJ-like decomposition Λ of A is a map whose restriction to each non surface type vertex group is a conjugation, and which sends surface type vertex groups on non-abelian images.

Proposition 8.1: *Let G be a torsion-free hyperbolic group, and let Λ a cyclic JSJ-like decomposition of G . Let $\theta : G \rightarrow G$ be a preretraction with respect to Λ , which sends surface type vertex groups of Λ isomorphically to conjugates of themselves. Then θ is an isomorphism.*

Proof. First note that if G is cyclic, the only JSJ-like decomposition it admits is the trivial one, for which the result is immediate. We may thus assume that G is not cyclic.

Denote by T the Bass-Serre tree T_Λ corresponding to Λ . To prove the proposition, we will construct a bijective simplicial map $j : T \rightarrow T$, such that j is equivariant with respect to θ in the following sense:

$$\forall g \in G, \forall v \in V(T), j(g \cdot v) = \theta(g) \cdot j(v).$$

For an edge e and a vertex v of T , the stabilisers of e and v in the standard action of G on T are denoted by G_e and G_v respectively.

1. Construction of the map j on vertices. By hypothesis, for each vertex v of T , there is an element g_v of G such that $\theta(G_v) = g_v G_v g_v^{-1}$. We set the image of v by j to be $g_v \cdot v$. Its stabiliser is exactly $\theta(G_v)$, and by remark 7.6, distinct vertices have distinct stabilisers, so this property defines $j(v)$ uniquely. Thus the image of $g \cdot v$ by j is the unique vertex whose stabiliser is $\theta(g)\theta(G_v)\theta(g^{-1})$, namely $\theta(g) \cdot j(v)$, and the map $v \mapsto j(v)$ is equivariant. Note that $j(v)$ is in the orbit of v , and thus is of the same type. Note also that $G_{j(v)} = \theta(G_v) \simeq G_v$.

2. The map $v \mapsto j(v)$ can be extended to a simplicial map $j : T \rightarrow T$. We need to check that adjacent vertices are sent on adjacent vertices. Suppose v, w adjacent, without loss of generality G_v is not a surface type vertex group. The intersection $G_v \cap G_w$ is an infinite cyclic group. On G_v , the map θ is just conjugation by the element g_v of G , so if we let $C := \theta(G_v \cap G_w)$, the group C is infinite cyclic. Moreover, C is contained in $\theta(G_v) \cap \theta(G_w)$. This means that

$j(v), j(w)$ are fixed by C , thus are at a distance at most 2. We will first show that it cannot be 2, then that it cannot be 0.

- Assume the distance is 2. The vertex u between $j(v)$ and $j(w)$ is a Z type vertex, which implies in particular that $j(v)$ and $j(w)$, and thus v and w , are not Z type vertex. Note that since $G_v \cap G_w$ is a subgroup of G_v , and since θ on G_v is just conjugation by g_v , we have $C = \theta(G_v \cap G_w) = g_v(G_v \cap G_w)g_v^{-1} < g_v G_w g_v^{-1}$, so it fixes the vertex $g_v \cdot w$. This vertex is at a distance 1 from $j(v)$, thus it is distinct from $j(v)$ and from $j(w)$. Its stabiliser is not cyclic, thus it is distinct from u . Hence we get a situation where C stabilises points $j(w)$ and $g \cdot w$ which are at a distance 3 of each other. This is a contradiction.
- Assume now $j(v) = j(w)$. Thus v and w are in the same orbit (in particular they must be of rigid type, since they are adjacent). Let $a \in G$ be such that $w = a \cdot v$. We have $G_w = aG_v a^{-1}$. We see that $\theta(a) \in \theta(G_v)$, since $j(v) = j(w) = j(a \cdot v) = \theta(a) \cdot j(v)$ and the stabiliser of $j(v)$ is $\theta(G_v)$. Thus there exists $a' \in G_v$ such that $\theta(a') = \theta(a)$.

Let $C_1 := G_v \cap G_w$, i.e. C_1 is the stabiliser of the edge e between v and w . Let $C_2 \leq G_v$ be such that $C_1 = aC_2a^{-1}$. Let c_1 generate C_1 , and $c_2 := a^{-1}c_1a$ generate C_2 . Note that by remark 7.7, C_1 is maximal abelian in G since it is the stabiliser of an edge which connects two rigid vertices. Now $\theta(c_2) = \theta(a^{-1})\theta(c_1)\theta(a)$ so that $\theta(c_2) = \theta(a'^{-1}c_1a')$. By injectivity of θ on G_v , $c_2 = a'^{-1}c_1a'$. Thus $a'a^{-1}$ commutes with c_1 , so it must be in C_1 and thus in G_v . But $a' \in G_v$ so we deduce $a \in G_v$ and $G_w = aG_v a^{-1} = G_v$. Since distinct vertices have distinct stabilisers, we get a contradiction.

Thus we can extend $v \mapsto j(v)$ to a simplicial map $j : T \rightarrow T$.

3. Injectivity of j . It is enough to show that there are no foldings, i.e. that no two edges adjacent to a same vertex are sent to the same edge by j . Suppose that two vertices w, w' of T are adjacent to a vertex v , and that the edges $e = [v, w]$ and $e' = [v, w']$ are sent on a same image by j . Let g_e be a generator of the stabiliser G_e of e , and $g_{e'}$ a generator of the stabiliser $G_{e'}$ of e' .

First it is clear that G_w and $G_{w'}$ must be conjugate since $j(w) = j(w')$, so w and w' are in the same orbit. Let $\gamma \in G$ such that $w' = \gamma \cdot w$. Note that $\gamma \notin G_w$.

Let us see that v must be a Z type vertex. We know that the stabiliser of $j(e)$ contains $\theta(g_e)$ and $\theta(g_{e'})$, so that $\theta([g_e, g_{e'}]) = 1$. As θ is injective on G_v , the elements g_e and $g_{e'}$ of G_v commute. Thus they have a common power which fixes both e and e' : by strong 2-acylindricity, v is a Z type vertex. This implies that w, w' , and $j(w)$ are not type Z vertices.

Remark that $\theta(\gamma) \cdot j(w) = j(\gamma \cdot w) = j(w') = j(w)$. Thus $\theta(\gamma)$ stabilises $j(w)$, hence it lies in $\theta(G_w)$. We can thus pick an element a of G_w such that $\theta(a) = \theta(\gamma)$.

Let g be an element of G_v which stabilise both e and e' : then g is both in G_w and in $\gamma G_w \gamma^{-1}$. Let $g' \in G_w$ be such that $g = \gamma g' \gamma^{-1}$. We have

$$\begin{aligned} \theta(g) &= \theta(\gamma)\theta(g')\theta(\gamma^{-1}) \\ &= \theta(a)\theta(g')\theta(a^{-1}) = \theta(ag'a^{-1}). \end{aligned}$$

Since θ is injective on G_w , we deduce that $g = ag'a^{-1}$ so $g' = \gamma^{-1}g\gamma = a^{-1}ga$. This shows $[\gamma a^{-1}, g] = 1$, so γa^{-1} preserves the set $\text{Fix}(g)$ of fixed point of g . But $\text{Fix}(g)$ has diameter 2 and is centred on v , so γa^{-1} fixes v , and $\gamma a^{-1} \in G_v$. Now a was chosen so that $\theta(\gamma) = \theta(a)$, so $\theta(\gamma a^{-1}) = 1$. By injectivity of θ on G_v , we get $\gamma = a$. This is a contradiction since γ is not in G_w , but a is.

4. Injectivity of θ . We have proved that j is injective, and this implies that θ is injective: if g is a non-trivial element of G , there exists $x \in T$ such that $g \cdot x \neq x$. Thus $j(g \cdot x) \neq j(x)$, so $\theta(g) \cdot j(x) \neq j(x)$ and $\theta(g)$ is non-trivial.

5. Surjectivity of j . We prove this by showing that if a vertex v is in the image of j , all the edges adjacent to v are also in the image. Suppose v is in the image of j , there exists g^v in G such that $j(g^v \cdot v) = v$. Pick e_1, \dots, e_r some representatives of the orbits of edges adjacent to v . The image e'_k of $g^v \cdot e_k$ by j must be adjacent to v .

We claim that if e_k and e_l lie in different orbits, so do e'_k and e'_l . Indeed, if e'_k and e'_l are in the same orbit, there exists α in G_v such that $\alpha \cdot e'_k = e'_l$. Since the action has no inversions, α must fix v . As v is in the image of j , its stabiliser is in the image of θ so there exists $a \in G$ such that $\theta(a) = \alpha$. Thus $\theta(a) \cdot j(g^v \cdot e_k) = j(g^v \cdot e_l)$, so by equivariance of j we get $j(ag^v \cdot e_k) = j(g^v \cdot e_l)$. By injectivity of j this means e_k and e_l are in the same orbit: this proves the claim. Thus the edges e'_k form a system of representative of the orbits of edges adjacent to v .

Now let e be an edge adjacent to v : there is an edge e'_k which is in the orbit of e , thus there is an element $\beta \in G$ such that $\beta \cdot e'_k = e$. Since the action has no inversions, β must fix v . We know G_v is in the image of θ so there exists $b \in G$ such that $\theta(b) = \beta$. Thus $j(b \cdot (g^v \cdot e_k)) = \theta(b) \cdot j(g^v \cdot e_k) = \beta \cdot e'_k = e$, so e is in the image of j . Hence all the vertices which neighbour v are in the image of j . This local surjectivity condition implies global surjectivity of j .

6. Surjectivity of θ . Let $g \in G$ and let v be a vertex of T with non-cyclic stabiliser. By surjectivity of j there exists w such that $j(w) = v$, and w' such that $j(w') = g \cdot v$. Clearly w and w' are in the same orbit. Thus there exists $h \in G$ such that $G_{w'} = hG_w h^{-1}$. We see that

$$gG_v g^{-1} = G_{g \cdot v} = \theta(G_{w'}) = \theta(h)\theta(G_w)\theta(h^{-1}) = \theta(h)G_v\theta(h^{-1}).$$

We get $G_v = g^{-1}\theta(h)G_v\theta(h)^{-1}g$. Thus G_v stabilises both v and $g^{-1}\theta(h) \cdot v$. Since G_v is not cyclic, $v = g^{-1}\theta(h) \cdot v$ so $g^{-1}\theta(h) \in G_v$. Since we know that G_v is in the image of θ , we get that g is in the image of θ .

We proved that θ is bijective, this terminates the proof. \square

Chapter 9

Non-pinching maps and the finite index property

In this chapter, we study morphisms $f : A \rightarrow G$, where G admits a decomposition as a graph of groups with surfaces Γ , first in the case where A is the fundamental group of a surface with boundary Σ , then more generally if A is the fundamental group of a graph of groups with surfaces Λ , or a free product of such groups. One of the aims is to give conditions under which any surface type vertex of Γ intersects the image of the morphism f either in a boundary subgroup, or in a subgroup of finite index. This is what we call the finite index property. One of the assumptions we will need is that the map f is non-pinching, that is, that its kernel does not contain elements corresponding to simple closed curves on Σ or on the surfaces of Λ . We will also see that under the right hypotheses, if a surface type vertex group S of Γ intersects $f(A)$ with finite index, it must contain the image of the fundamental group of a subsurface of a surface Σ' of Λ . We will also show that this implies that the complexity of the surface Σ' is greater than that of the surface Σ corresponding to S .

9.1 Surfaces with boundary

We first restrict ourselves to the case where A is the fundamental group of a surface with boundary.

9.1.1 Surfaces acting on simplicial trees

Let us first give a useful lemma to understand actions of fundamental groups of surfaces with boundary on simplicial trees. For this, we need the following definitions.

Definition 9.1: (essential curves, elements corresponding to an essential curve, essential tubular neighbourhood) *An essential curve γ on a surface with boundary Σ is the free homotopy class of a non-contractible, two-sided, and non-boundary parallel simple closed curve γ_0 .*

To an essential curve γ corresponds a conjugacy class of infinite cyclic subgroups of the fundamental group of Σ , we call their generators the elements corresponding to γ .

The simple closed curve γ_0 has an open neighbourhood which is homeomorphic to an annulus, we call such a neighbourhood a tubular neighbourhood of γ . Given a set of essential curves \mathcal{C} , a tubular neighbourhood \mathcal{C}_a of \mathcal{C} is the union of disjoint tubular neighbourhoods, one for each essential curve in \mathcal{C} .

Definition 9.2: (graph of groups $\Delta(\Sigma, \mathcal{C})$ dual to a set of curves) *Let Σ be a surface with boundary, and let \mathcal{C} be a set of essential curves on Σ . By the Van Kampen lemma, S admits a splitting*

$\Delta(\Sigma, \mathcal{C})$ whose edge groups are the infinite cyclic groups generated by elements corresponding to the curves of \mathcal{C} , and whose vertex groups are the fundamental groups of the connected components of the complement in Σ of a tubular neighbourhood \mathcal{C}_a of \mathcal{C} . An edge corresponding to a curve γ joins two vertices corresponding to connected components Σ_1 and Σ_2 if there is a path in Σ between Σ_1 and Σ_2 which intersects only one component of \mathcal{C}_a , the one corresponding to γ . We call $\Delta(\Sigma, \mathcal{C})$ the graph dual to the set of curves \mathcal{C} , and the corresponding tree $T_{\mathcal{C}}$ is called the tree dual to \mathcal{C} .

Finally we define

Definition 9.3: (minimal equivariant map) *Let G be a group which acts on simplicial trees T and T' . An equivariant map $t : T \rightarrow T'$ is said to be minimal if it sends vertices on vertices, if every edge is sent on the unique path between the images of its endpoints, and if for any vertex v of T whose stabiliser also stabilises an edge e adjacent to $t(v)$, no open neighbourhood of v has image contained in e .*

The following lemma is a particular case of theorem III.2.6 in [MS84].

Lemma 9.4: *Suppose that the fundamental group S of a surface with boundary Σ acts on a simplicial tree T , in such a way that the boundary subgroups are elliptic. Then there exists a system \mathcal{C} of essential curves on Σ , and a minimal equivariant map $t : T_{\mathcal{C}} \rightarrow T$.*

Note that the map t is not necessarily simplicial.

Remark 9.5: *The cyclic subgroups of S corresponding to curves in \mathcal{C} stabilise edges of T . If \mathcal{C}_a is a tubular neighbourhood of \mathcal{C} , the fundamental groups of connected components of the complement of \mathcal{C}_a are vertex groups of $\Delta(\Sigma, \mathcal{C})$, thus they are elliptic in T .*

The proof of this lemma is essentially the first part of the proof of theorem III.2.6 in [MS84]. Since we do not claim that the equivariant map is injective, we do not need to assume that the stabilisers in S of the edges of T are cyclic.

The idea of the proof is to construct an equivariant simplicial map f from a universal cover of Σ to the tree T , then to look at the lift by f of midpoints of edges of T . The map can be built in such a way that the lifts by f give by the covering map non-null homotopic simple closed curves. We take for \mathcal{C} the homotopy classes of these simple closed curves. Since some of the curves might be in the same homotopy class, we lose the simpliciality of the map f .

9.1.2 Non-pinching maps and the finite index property

Definition 9.6: (non-pinching) *Let Σ be a surface with boundary, and let S be its fundamental group. A morphism $f : S \rightarrow G$ is said to be non-pinching with respect to Σ if its kernel does not contain any element corresponding to an essential curve lying on Σ , and if it is injective on boundary subgroups.*

The following lemma is a crucial ingredient of the proof of proposition 7.15.

Lemma 9.7: *Let S and S' be the fundamental groups of surfaces with boundary Σ and Σ' . Let $f : S \rightarrow S'$ be a non-pinching map which sends boundary subgroups of S into boundary subgroups of S' . If $f(S)$ is not contained in a boundary subgroup of S' , then it is a subgroup of finite index of S' .*

To prove it, we will use

Lemma 9.8: *Let Q be the fundamental group of a surface with boundary Ξ . If Q_0 is a finitely generated infinite index subgroup of Q , it is of the form*

$$C_1 * \dots * C_m * F$$

where F is a (possibly trivial) free group, each of the groups C_j is a boundary subgroup of Q , and any boundary element of Q contained in Q_0 can be conjugated in one of the groups C_j by an element of Q_0 .

Proof. By Theorem 2.1 in [Sco78], there exists a finite covering $p : \Xi_1 \rightarrow \Xi$, and a subsurface Ξ_0 of Ξ_1 , such that Q_0 is the image by the injection p_* of the fundamental group of Ξ_0 . Let $Q_1 = \pi_1(\Xi_1)$. The covering is finite, so Ξ_1 is compact, Q_1 is of finite index in Q , and the boundary elements of Q_1 are exactly the boundary elements of Q contained in Q_1 . Since Q_0 is of infinite index in Q , it must be of infinite index in Q_1 . Thus Ξ_0 is a proper subsurface of Ξ_1 , and thus at least one of its boundary components γ is not a boundary component of Ξ_1 . This implies the lemma. \square

We can now prove lemma 9.7.

Proof. Suppose $f(S)$ has infinite index in S' . Then it admits a free product decomposition $C_1 * \dots * C_m * F$ as given by lemma 9.8, in which $m \geq 1$ since boundary elements of S are sent to boundary elements of S' . Since $f(S)$ is not contained in a boundary subgroup of S , this decomposition contains at least two factors, so the corresponding minimal $f(S)$ -tree T_0 with trivial edge stabilisers is not reduced to a point. The group S acts via f on T_0 , the tree T_0 is minimal for this action, and boundary subgroups of S are sent to boundary subgroups of S' , thus they lie in conjugates of the factors C_i and they are elliptic in T_0 . By lemma 9.4, we get a set of essential simple closed curves on Σ whose corresponding elements stabilise edges of T_0 via f , i.e. have trivial image by f . This contradicts the fact that f is non-pinching on Σ . \square

9.1.3 Complexities

We will denote by $\text{rk}(F)$ the rank of a finitely generated free group F .

Definition 9.9: (topological complexity) *Let Σ be a surface with boundary, denote by S its fundamental group. The topological complexity $k(\Sigma)$ of Σ is the pair $(\text{rk}(S), -n)$, where n is the number of boundary components of Σ . We order topological complexities by the lexicographic order.*

We will give a lemma which shows in particular that if we have a non-pinching morphism as above between the fundamental groups of surfaces with boundary Σ and Σ' , then the complexity of Σ is at least the complexity of Σ' .

Lemma 9.10: *Let S and S' be the fundamental groups of surfaces with boundary Σ and Σ' . If $f : S \rightarrow S'$ is a map which sends boundary subgroups of S into boundary subgroups of S' , and such that $f(S)$ is a subgroup of finite index of S' , then*

$$k(\Sigma) \geq k(\Sigma');$$

and if we have equality, f is an isomorphism.

Proof. A subgroup of finite index in a finitely generated free group of rank n is a free group of rank at least n , with equality if and only if the index is 1. Thus $\text{rk}(S') \leq \text{rk}(f(S))$ with equality if and only if f is surjective. Now $\text{rk}(f(S)) \leq \text{rk}(S)$, and since free groups are Hopfian, we have equality if and only if f is injective. Thus $\text{rk}(S') \leq \text{rk}(S)$, with equality if and only if f is an isomorphism. If this is the case, f sends non-conjugate boundary subgroups of S to non-conjugate boundary subgroups of S' , so that Σ' has at least as many boundary component as Σ . Thus $k(\Sigma') \leq k(\Sigma)$. \square

9.2 Graphs of groups with surfaces

We now want to generalise the previous sections from the case where A is fundamental group of a surface with boundary to the case where A is the fundamental group of a graph of groups with surfaces.

9.2.1 Elliptic refinements of graphs of groups with surfaces

Let A and G be fundamental groups of graphs of groups with surfaces Λ and Γ respectively. Let $f : A \rightarrow G$ be a morphism which sends edge groups and non surface type vertex groups of Λ into non surface type vertex groups of Γ .

Each surface type vertex group S of Λ corresponding to a surface Σ acts on the tree T_Γ corresponding to Γ via the map f , and boundary subgroups of S are elliptic in this action. By lemma 9.4, we get a set of essential curves $\mathcal{C}^+(\Sigma)$ on Σ . We can then refine the graph of groups Λ by the graph of groups $\Delta(\Sigma, \mathcal{C}^+(\Sigma))$ dual to the set of curves $\mathcal{C}^+(\Sigma)$ (recall definition 9.2). Every vertex group of the refined graph of groups Λ^+ thus obtained is elliptic in the action of A on T_Γ via f . Denote by \mathcal{C}^+ the union of all the sets $\mathcal{C}^+(\Sigma)$.

Definition 9.11: (elliptic refinement of a graph of group) *We call the graph of groups Λ^+ built as above an elliptic refinement of Λ relative to f and Γ , given by the set of curves \mathcal{C}^+ .*

Remark 9.12: *There is a map $t^+ : T_{\Lambda^+} \rightarrow T_\Gamma$ which sends vertices on vertices, is f -equivariant and minimal. This an easy consequence of the fact that all the vertex groups of Λ^+ have image by f elliptic in T_Γ .*

9.2.2 Non-pinching maps on graphs of groups with surfaces

Definition 9.13: (non-pinching with respect to a graph of groups with surfaces) *Let Λ be a graph of groups with surfaces. We say that a morphism $f : \pi_1(\Lambda) \rightarrow G$ is non-pinching with respect to Λ if the restriction of f to each surface type vertex group of Λ is non-pinching.*

Setting. For the rest of section 9.2, A_1 and A are groups which admit decompositions Λ_1 and Λ as a graph of groups with surfaces whose edge groups are infinite cyclic. Also, $f : A_1 \rightarrow A$ is a morphism which sends non surface type vertex groups and edge groups of Λ_1 injectively into non surface type vertex groups and edge groups of Λ respectively.

By the previous section, we can define an elliptic refinement Λ_1^+ of Λ_1 with respect to f and Λ . We then know that there exists a minimal f -equivariant map $t^+ : T_{\Lambda_1^+} \rightarrow T_\Lambda$. In the case where f is non-pinching with respect to Λ_1 , the next lemma gives us necessary and sufficient conditions on a surface type vertex of T_Λ for it to lie in the image of t^+ .

Lemma 9.14: *Suppose we are in the setting above. If f is non-pinching with respect to Λ_1 , for any surface type vertex v of T_Λ with stabiliser S the following are equivalent*

- (i) v lies in the image of $T_{\Lambda_1^+}$ by t^+ ;
- (ii) there is a conjugate of a surface type vertex group S^+ of Λ_1^+ whose image by f lies in S ;
- (iii) there is a conjugate of a surface type vertex group S^+ of Λ_1^+ whose image by f is a subgroup of finite index of S ;
- (iv) the intersection of S with the image of A_1 by f is not contained in a boundary subgroup of S .

Proof. The fact that f is non-pinching on Λ_1 , and injective on its edge groups implies that it is also non-pinching on Λ_1^+ and injective on its edge groups.

(i) \Rightarrow (ii): If w is a non surface type vertex of $T_{\Lambda_1^+}$ with non-abelian stabiliser R , then f is injective on R so $f(R)$ is non-abelian, thus it stabilises exactly one vertex in T_Λ . But $f(R)$ lies in a non surface type vertex group of Λ , so $t^+(w) \neq v$.

Suppose now that w is a non surface type vertex w of $T_{\Lambda_1^+}$ with abelian stabiliser Z . By minimality, either the image of the star of w intersects at least two edges adjacent to $t^+(w)$, or $f(Z)$ properly contains all the edge group of the unique edge on which the star of w is sent. In the first case, note that the image of the star of w is stabilised by a non-trivial element, so by 1-acylindricity next to surface type vertices, $t^+(w)$ is not of surface type. In the second case, note that edge groups adjacent to surface type vertices are maximal cyclic in the surface group, so $t^+(w)$ can not be of surface type.

Finally, if e is an edge of $T_{\Lambda_1^+}$, the image of its interior is stabilised by a non-trivial element, thus it does not contain any surface type vertices by 1-acylindricity next to surface type vertices and by minimality of t^+ .

Thus we see that if a surface type vertex with stabiliser S is in the image of t^+ , it means that it is the image of some surface type vertex of $T_{\Lambda_1^+}$ with stabiliser S^+ . Thus $f(S^+) \leq S$ as claimed.

(ii) \Rightarrow (iii): The map f sends edge groups of $T_{\Lambda_1^+}$ to edge groups of T_Λ , thus boundary subgroups of S^+ are sent to boundary subgroups of S . Moreover, by minimality of t^+ and 1-acylindricity next to surface type vertices, $f(S^+)$ is not contained in a boundary subgroup of S . By lemma 9.7, this means that $f(S^+)$ has finite index in S .

(iii) \Rightarrow (iv): This is clear.

(iv) \Rightarrow (i): If v lies outside of $t^+(T_{\Lambda_1^+})$, the intersection between $f(A_1)$ and S stabilises both v and $t^+(T_{\Lambda_1^+})$, thus it stabilises the non-trivial path between them. Thus it stabilises one of the edges adjacent to v , which implies that it is contained in a boundary subgroup of S . \square

9.2.3 Surface complexity of graphs of groups

Definition 9.15: (complexity of a set of surfaces, surface complexity of a graph of groups with surfaces) Let $\mathcal{S} = \{\Sigma_i \mid 1 \leq i \leq l\}$ be a set of surfaces with boundary, and recall that $k(\Sigma_i)$ denotes the topological complexity of Σ_i . The complexity $K(\mathcal{S})$ is the finite sequence $(k(\Sigma_i))_{1 \leq i \leq l}$ of the complexities of surfaces of \mathcal{S} arranged in decreasing order.

We order the complexities of sets of surfaces lexicographically, that is

$$k(\Sigma_1) \dots k(\Sigma_l) < k(\Sigma'_1) \dots k(\Sigma'_{l'})$$

if $\{i \mid k(\Sigma_i) \neq k(\Sigma'_i); 1 \leq i \leq \min\{l, l'\}\}$ is non-empty, has minimum j , and $k(\Sigma_j) < k(\Sigma'_j)$; or if the set is empty and $l < l'$.

If Λ is a graph of groups with surfaces, its surface complexity is the complexity of its set of surfaces.

Lemma 9.16: If \mathcal{C}^+ is not empty, the surface complexity of an elliptic refinement Λ_1^+ of a graph of groups Λ_1 is strictly smaller than that of Λ_1 .

Proof. Let Σ be a surface of Λ_1 with fundamental group S . The vertex corresponding to Σ in Λ_1 is replaced by the graph of groups $\Delta(\Sigma, \mathcal{C}^+)$ to build Λ_1^+ . Thus, showing that the surfaces of $\Delta(\Sigma, \mathcal{C}^+)$ have complexity strictly smaller than that of Σ is enough to prove the lemma.

The rank of S is given by $1 - \chi(\Sigma)$, where χ is the Euler characteristic. Suppose \mathcal{C}^+ contains a single curve which lies on Σ . If it separates Σ into two subsurfaces Σ_1 and Σ_2 , we have

$$\chi(\Sigma) = \chi(\Sigma_1) + \chi(\Sigma_2)$$

Since the curves are not boundary parallel, Σ_1 and Σ_2 have strictly negative Euler characteristic, or one of the subsurfaces, without loss of generality Σ_2 , is a Möbius band. In the first case, the characteristic of Σ_1 and Σ_2 are strictly bigger than that of Σ . In the second case, Σ_1 has the same Euler characteristic as Σ , but one extra boundary component. Thus the complexities of Σ_1 and Σ_2 are always strictly smaller than that of Σ . If the curve is not separating, $\Delta(\Sigma, \mathcal{C}^+)$ has a unique surface type vertex, whose corresponding surface has Euler characteristic equal to that of Σ , but which has two additional boundary components. Its complexity is therefore strictly smaller than that of Σ . If more than one curve lies on Σ , we proceed by induction. \square

The following lemma gives us a relation between the surface complexities of graphs of groups Λ_1 and Λ when the map $f : A_1 \rightarrow A$ is non-pinching.

Lemma 9.17: *If f is non-pinching with respect to Λ_1 , and if $t^+ : T_{\Lambda_1^+} \rightarrow T_\Lambda$ is surjective, the surface complexity of Λ_1 is greater than or equal to that of Λ .*

Proof. Each surface type vertex lies in the image of t^+ , so by lemma 9.14, for each surface type vertex group S of Λ there is a surface type vertex group S^+ of Λ_1^+ such that $f(S^+)$ has finite index in a conjugate of S . By lemma 9.10, the complexity of the surface corresponding to S^+ is thus greater than that of the surface corresponding to S . In this way, to each surface of Λ corresponds a surface of Λ_1^+ whose complexity is greater, and this correspondence gives an injection from the set of surfaces of Λ to the set of surfaces of Λ_1^+ . This implies that the surface complexity of Λ is smaller than that of Λ_1^+ , which in turn is smaller than that of Λ_1 by lemma 9.16. \square

9.3 Finite index property for free products

We now want to prove a proposition that should be thought of as a generalisation of lemma 9.14 in the case where instead of a morphism from A_1 to A , we have a map from a free product $A_1 * \dots * A_l$ to A . We will see that up to conjugation on these free factors, we still control which surface type vertex groups of Λ intersect the image of a non-pinching map in a subgroup bigger than a boundary subgroup.

Proposition 9.18: *Let A_1, \dots, A_l be groups which admit JSJ-like decompositions $\Lambda_1, \dots, \Lambda_l$ and let Λ be a graph of groups with surfaces with fundamental group A . Assume that $K(\Lambda_i) < K(\Lambda)$.*

*Suppose $h : A_1 * \dots * A_l \rightarrow A$ is a map which sends non surface type vertex groups and edge groups of the graphs of groups Λ_i injectively into non surface type vertex groups and edge groups of Λ respectively, and such that the maps $h|_{A_i}$ are non-pinching with respect to the graphs Λ_i . For each i with $1 \leq i \leq l$, let Λ_i^+ be an elliptic refinement of Λ_i with respect to $h|_{A_i}$ and Λ .*

*Then there exists a map $\tilde{h} : A_1 * \dots * A_l \rightarrow A$ such that $\tilde{h}|_{A_i}$ coincides with $h|_{A_i}$ up to conjugation, such that $\tilde{h}(A_1 * \dots * A_l) = \tilde{h}(A_1) * \dots * \tilde{h}(A_l)$, and such that for any surface type vertex group S of Λ , the following are equivalent:*

- (i) *The intersection of S with $\tilde{h}(A_1 * \dots * A_l)$ is not contained in a boundary subgroup of S .*
- (ii) *There is a conjugate of a surface type vertex group S^+ of one of the graphs of groups Λ_i^+ whose image by \tilde{h} has finite index in S .*

To prove this we will need the following lemmas.

Lemma 9.19: *Let G be a finitely generated group, and let T be a minimal irreducible G -tree. If τ and τ' are proper subtrees of T , for any integer D , there is a translate of τ' by an element of G which lies at a distance at least D of τ .*

Proof. By lemme 4.3 in [Pau89], the hypotheses allow us, for any two distinct vertices v and w of T , to find an element of G which is hyperbolic in the action of G on T , and whose axis contains the path between v and w .

Suppose first that the smallest tree τ_0 containing $\tau \cup \tau'$ is a proper subtree of T . Let K be a connected component of the complement of τ_0 in T , and let u be the vertex of T such that $\overline{K} \cap \tau_0 = \{u\}$. By minimality and irreducibility of the action, K is not a line, so we can find points v and w in such a component such that the tripod formed by v , w , and u is non-trivial. We pick a hyperbolic element g whose axis contains the path between v and w . The projection of τ and τ' on the axis of g is reduced to a point. Thus $g^D \cdot \tau'$ is at distance greater than D of τ .

If on the other hand, $\tau_0 = T$, we pick vertices v, w of the tree which are in τ' but not in τ , and in τ but not in τ' respectively. Now τ lies in the connected component of $T - \{v\}$ containing w and τ' lies in the connected component of $T - \{w\}$ containing v . Thus the intersection $\tau \cap \tau'$ lies in the connected component of $T - \{v, w\}$ containing the arc between v and w . Pick a hyperbolic element whose axis contains the path between v and w . By applying a suitable power of this element we can translate τ' away from τ . \square

Lemma 9.20: *Let G be a finitely generated group, and let τ be a k -acylindrical minimal G -tree. Suppose G_1 and G_2 are subgroups of G which generate G , and whose minimal subtrees T_1 and T_2 in τ lie at a distance at least $2k + 3$ from each other. If v is a vertex which lies in τ*

- *either $\text{Stab}(v)$ stabilises an edge adjacent to v ;*
- *or v lies in a translate of T_i by an element of G , and in this case $\text{Stab}(v)$ stabilises this translate.*

Remark 9.21: *If the hypotheses hold, we have $G = G_1 * G_2$. Indeed, the minimal tree τ of G is the union of translates of T_1 , translates of T_2 , and translates of the path between them. Since the path between them has length greater than $k + 1$, it is trivially stabilised.*

Proof. Denote by D the path joining T_1 to T_2 . The tree τ is the union of translates of T_1 , T_2 and D by elements of G . Let \hat{T}_i for $i = 1, 2$ be the set of points whose distance to T_i is at most $k + 1$: note that \hat{T}_1 and \hat{T}_2 are disjoint. Denote by \hat{D} the subsegment of D which joins \hat{T}_1 and \hat{T}_2 . Let B_1 be the complement in $\tau - \hat{T}_1$ of the connected component containing the interior of \hat{D} , and let B_2 be the complement in $\tau - \hat{T}_2$ of the connected component containing the interior of \hat{D} .

By k -acylindricity, an element of G_1 sends points of \hat{D} , of \hat{T}_2 and of B_2 into B_1 , and an element of G_2 sends points of \hat{D} , of \hat{T}_1 and of B_1 into B_2 .

If $v \in \hat{D}$, its image by a non-trivial element of G lies in $B_1 \cup B_2$, thus $\text{Stab}(v) \cap G$ is trivial. If $v \in \hat{T}_1$ and $g \cdot v = v$ then $g \in G_1$: indeed, otherwise we can see that $g \cdot v \in B_2$. Thus if $v \in T_1$, the stabiliser of v also stabilises T_1 , and if $v \in \hat{T}_1 - T_1$, the stabiliser of v also stabilises the path between v and T_1 , so it stabilises an edge adjacent to v . We get a similar result if $v \in \hat{T}_2$. If v lies in a translate $g \cdot \hat{D}$ of \hat{D} , or in a translate $g \cdot \hat{T}_i$ of \hat{T}_i , we apply the results above to $g^{-1} \cdot v$.

This is enough to conclude. \square

We can now prove proposition 9.18.

Proof. We prove by induction on the number of factors l that the result holds, and that moreover we can require that the map \tilde{h} is such that the minimal subtree of $\tilde{h}(A_1 * \dots * A_l)$ in T_Λ is a proper subtree.

For $l = 1$, if we take $\tilde{h} = h$ the result holds by 9.14. Since we assumed that $K(\Lambda_1) < K(\Lambda)$, the minimal subtree of $h(A_1)$ does not cover T_Λ by lemma 9.17.

Suppose by induction that for $l = n - 1$, the induction hypothesis holds. Let h be a map $A_1 * \dots * A_n \rightarrow A$ which satisfies all the hypotheses. The induction hypothesis gives us a map

\tilde{h} from $A_1 * \dots * A_{n-1}$ to A such that $\tilde{h}|_{A_i}$ coincides with $h|_{A_i}$ up to conjugation for $i < n$, and such that the minimal subtree T_1 of $G_1 = \tilde{h}(A_1 * \dots * A_{n-1})$ is a proper subtree of T_Λ .

Consider the minimal tree of $h(A_n)$: since we assumed $K(\Lambda_n) < K(\Lambda)$, by lemma 9.17, it is also a proper subtree of T_Λ . Thus by lemma 9.19, it has a translate T_2 by an element α of A which lies at a distance at least 7 of T_1 . Extend \tilde{h} to A_n by setting $\tilde{h}|_{A_n} = \text{Conj}(\alpha) \circ h|_{A_n}$. Then T_2 is the minimal subtree of $G_2 = \tilde{h}(A_n)$.

Note that the group G generated by $G_1 = \tilde{h}(A_1 * \dots * A_{n-1})$ and $G_2 = \tilde{h}(A_n)$ is precisely $\tilde{h}(A_1 * \dots * A_n)$, we denote its minimal subtree by τ and we apply lemma 9.20. By remark 9.21, $\tilde{h}(A_1 * \dots * A_n) = \tilde{h}(A_1 * \dots * A_{n-1}) * \tilde{h}(A_n)$ so by induction hypothesis we get

$$\tilde{h}(A_1 * \dots * A_n) = \tilde{h}(A_1) * \dots * \tilde{h}(A_l).$$

Moreover, τ is properly contained in T_Λ , since the points which lie on the path between T_1 and T_2 are branching points in T_Λ , but not in τ .

Now let v be a surface type vertex of T_Λ , and denote by S its stabiliser. If v lies outside of τ , the intersection $S \cap \tilde{h}(A_1 * \dots * A_n)$ stabilises both v and τ , thus it is contained in a boundary subgroup of S . We may thus assume that v lies in τ . By lemma 9.20, either $S \cap \tilde{h}(A_1 * \dots * A_n)$ is contained in the stabiliser of an edge adjacent to v , in which case we are done, or v lies in a translate of T_1 or T_2 .

If v lies in T_1 itself, lemma 9.20 also tells us that the stabiliser of v by G , i.e. the intersection $S \cap \tilde{h}(A_1 * \dots * A_n)$, is contained in the stabiliser G_1 of T_1 , namely $\tilde{h}(A_1 * \dots * A_{n-1})$. By induction hypothesis we have two possibilities: either the intersection $S \cap \tilde{h}(A_1 * \dots * A_{n-1})$ lies in a boundary subgroup of S , but then so does the intersection $S \cap \tilde{h}(A_1 * \dots * A_n)$; or there is a conjugate of a surface type vertex group S^+ of one of the graphs Λ_i^+ for $i < n-1$ whose image by \tilde{h} lies in the stabiliser of v .

If v lies in T_2 itself, lemma 9.20 also tells us that the intersection $S \cap \tilde{h}(A_1 * \dots * A_n)$ is contained in $\tilde{h}(A_n)$. Then, by lemma 9.14, there is a conjugate of surface type vertex group S^+ of Λ_n^+ whose image by \tilde{h} lies in the stabiliser of v .

Finally, if v lies in a translate of T_1 or T_2 by an element α of $\tilde{h}(A_1 * \dots * A_n)$, we apply the results above to the vertex $\alpha^{-1} \cdot v$. This is enough to prove the result. \square

Chapter 10

From preretractions to hyperbolic floors

In this chapter, we prove proposition 7.15 and proposition 7.16. From the existence of a non-injective preretraction $f : A \rightarrow A$, proposition 7.15 deduces the existence of a retraction r which makes $(A, r(A), r)$ a hyperbolic floor, and from the existence of a non-injective preretraction $A \rightarrow G$, proposition 7.16 deduces the existence of a preretraction from A to a retract of G . In both proofs, the idea is to modify f into the retraction r .

The previous chapter showed that for a non-pinching map, we control what happens to surface type vertices: it will thus be useful to work with non-pinching maps. The first section of this chapter explains how to factor a preretraction $f : A \rightarrow G$ as $f' \circ \rho$, where f' is non-pinching with respect to some free factors of $\rho(A)$. This will be done by letting ρ kill elements corresponding to simple closed curves which lie in the kernel of f .

In the second section, we worry about the non-abelianity of the image of surface groups, and give a criterion which will enable us later to guarantee that it is still satisfied despite all the transformations we will make f undergo.

In the third section, we define a complexity on the set of non-injective preretractions $A \rightarrow A$. Note that this set contains f , so is non-empty by hypothesis. We then proceed to study a maximal element, and we will see how we can build from it a retraction $A \rightarrow A'$ which makes (A, A', r) a hyperbolic floor, thus proving proposition 7.15.

The fourth section finally gives a proof of proposition 7.16. It should be noted that the third and the fourth section are independent.


10.1 Pinching a set of curves

Let A be the fundamental group of a graph of group with surfaces Λ which has infinite cyclic edge groups. Let \mathcal{C} be a set of essential curves on the surfaces of Λ . Let $N(\mathcal{C})$ be the subgroup of A normally generated by the elements corresponding to the curves of \mathcal{C} .

Definition 10.1: (pinching map) *We denote by $\rho_{\mathcal{C}}$ the quotient map $A \rightarrow A/N(\mathcal{C})$, and we call it the pinching map of A by \mathcal{C} .*

Denote by $\rho_{\mathcal{C}}$ of Λ the quotient decomposition, namely the decomposition obtained from Λ by replacing each vertex group by its image by $\rho_{\mathcal{C}}$ (not that $\rho_{\mathcal{C}}$ is injective on edge groups of Λ).

Let us now build a decomposition as a graph of groups with surfaces for $\rho_{\mathcal{C}}(A)$. For this, we will refine $\rho_{\mathcal{C}}(\Lambda)$ by decompositions of the groups $\rho_{\mathcal{C}}(S)$. For each surface type vertex group S of Λ with corresponding surface Σ , consider the graph of groups dual to the set of essential curves

Figure 10.1: The pinching of Λ by \mathcal{C} .

of \mathcal{C} which lie on Σ . We denote it by $\Delta(\Sigma, \mathcal{C})$. We get a graph of group decomposition $\Gamma(\Sigma, \mathcal{C})$ for $\rho_{\mathcal{C}}(S)$ by replacing each vertex and edge group of $\Delta(\Sigma, \mathcal{C})$ by its image by $\rho_{\mathcal{C}}$.

A vertex group S_0 of $\Delta(\Sigma, \mathcal{C})$ is the fundamental group of a subsurface Σ_0 of Σ . The image of S_0 by $\rho_{\mathcal{C}}$ is the fundamental group of the surface obtained by gluing discs to the boundary components of Σ_0 corresponding to curves of \mathcal{C} . Note thus that if all the boundary components of Σ_0 correspond to curves of \mathcal{C} , the image of S_0 by $\rho_{\mathcal{C}}$ is the fundamental group of a closed surface. Then, we call the corresponding vertex of $\Gamma(\Sigma, \mathcal{C})$ an interior vertex.

Refine the graph of groups $\rho_{\mathcal{C}}(\Lambda)$ by replacing each surface type vertex with corresponding surface Σ by the graph of groups $\Gamma(\Sigma, \mathcal{C})$ (see figure 10.1).

Definition 10.2: (pinching of a graph of groups) *We call the graph of groups $\Lambda_{\mathcal{C}}$ thus obtained the pinching of Λ by \mathcal{C} .*

Let us see that this graph of groups decomposition gives us a decomposition of $\rho_{\mathcal{C}}(A)$ as a free product. Remove from $\Lambda_{\mathcal{C}}$ all the interior of edges of the graphs $\Gamma(\Sigma, \mathcal{C})$ as well as the interior vertices: denote by $\Lambda_1, \dots, \Lambda_l$ the various connected components. They are subgraphs of groups of $\Lambda_{\mathcal{C}}$, and they admit a natural structure of graph of groups with surfaces whose surface type vertices are exactly the vertices which belong to one of the subgraphs $\Gamma(\Sigma, \mathcal{C})$. Call A_1, \dots, A_l their fundamental groups.

Remark 10.3: *The graphs of groups Λ_i are JSJ-like decompositions.*

Lemma 10.4: *If \mathcal{C} is not empty, the complexity of the set containing all the surfaces of the graphs of groups Λ_i is strictly smaller than the complexity of the set of surfaces of Λ .*

Proof. As in the proof of lemma 9.16, it is enough to see that in the graph of groups $\Gamma(\Sigma, \mathcal{C})$ which replaces the vertex corresponding to Σ , all the surfaces have complexity smaller than that of Σ . But the surfaces of $\Gamma(\Sigma, \mathcal{C})$ are obtained from surfaces of $\Delta(\Sigma, \mathcal{C})$ by gluing discs to boundary components, which strictly decreases the Euler characteristic, and thus the complexity. We saw in the proof of 9.16 that if at least one curve of \mathcal{C} lies on Σ , the surfaces of $\Delta(\Sigma, \mathcal{C})$ have complexity strictly smaller than that of Σ . This terminates the proof. \square

Remark that if we collapse the edges of the subgraphs Λ_i in Λ , the graph of groups we get has trivial edges stabilisers. Picking a maximal subtree in it, and choosing a lift in the $\rho_C(A)$ -tree corresponding to this graph of group gives us an identification of the groups A_i to subgroups of $\rho_C(A)$, and a free product decomposition of $\rho_C(A)$ of the form

$$\rho_C(A) = (A_1 * \dots * A_l) * (S_1 * \dots * S_p) * (Z_1 * \dots * Z_q) \quad (\dagger)$$

where the groups S_j are fundamental groups of closed surfaces which are not spheres, corresponding to interior vertices of the graphs of groups $\Gamma(\Sigma, \mathcal{C})$, and each group Z_k is the infinite cyclic subgroup of $\rho_C(A)$ corresponding to an edge lying outside the maximal subtree.

Definition 10.5: (pinching decomposition of $\rho_C(A)$) *We call the free product decomposition (\dagger) a pinching decomposition of $\rho_C(A)$ with respect to \mathcal{C} .*

Note that different choices of maximal subtree and different lifts in the $\rho_C(A)$ -tree give us different pinching decompositions of $\rho_C(A)$.

Finally, we will use the following notations

$$\begin{aligned} A_C &:= A_1 * \dots * A_l \\ R_C &:= (S_1 * \dots * S_p) * (Z_1 * \dots * Z_q). \end{aligned}$$

Some of the vertex groups of the graphs of groups Λ_i which have been scarcely modified by the map ρ_C will play a particular role in the third section. They are given by

Definition 10.6: (intact surface type vertex of Λ_i) *Let Σ be a surface of Λ , and let S be the corresponding vertex group. Suppose that the graph of groups $\Gamma(\Sigma, \mathcal{C})$ is a tree of groups, all of whose vertex groups except one are trivial or isomorphic to $\mathbb{Z}/2\mathbb{Z}$. The exceptional vertex group S_0 is conjugate to a surface type vertex group of one of the graphs of groups Λ_i . We call such a surface type vertex of Λ_i an intact surface type vertex, and the corresponding surface is called an intact surface of Λ_i .*

We said that our strategy was to factor the non-injective preretraction f as $f = f' \circ \rho$, where the map f' is non-pinching with respect to a suitable graph of groups. The map ρ should thus be the quotient of A by a maximal set of elements coming from simple closed curves killed by f . Note that if f is injective on edge groups, no element corresponding to a boundary parallel simple closed curve lies in the kernel of f . To make this precise, we give

Definition 10.7: (essential curves killed by f) *Let A be the fundamental group of a graph of groups with surfaces Λ , whose edge groups are infinite cyclic. Let $f : A \rightarrow G$ be a map which is injective on edge groups.*

Consider systems of two-sided non-homotopic non boundary parallel simple closed curves on the surfaces of Λ whose corresponding elements in A are in the kernel of f . For each curve in such a set, we say that the corresponding essential curve is killed by f .

If the system of simple closed curves we chose is maximal for inclusion among all such systems, the associated set of free homotopy classes is called a maximal set of essential curves of Λ killed by f .

Remark 10.8: *In the setting of definition 10.7, if \mathcal{C} is a maximal set of essential curves killed by f , the map f factors as $f' \circ \rho_C$, and $f'|_{A_i}$ is non-pinching with respect to Λ_i .*

10.2 Non-abelianity of surfaces

This criterion will prove very useful in the proofs of propositions 7.15 and 7.16. It will imply that if we have a map g from A_C to a torsion-free hyperbolic group G , as long as intact surface type

vertex groups are not sent to abelian images, we can extend g to a map from $\rho_C(A) = A_C * R_C$ to G whose composition with ρ_C sends all the surface type vertex groups of Λ to non-abelian images.

Lemma 10.9: *Let A be a group which admits a JSJ-like decomposition Λ . Let C be a set of essential curves on the surfaces of Λ . Choose a pinching decomposition of $\rho_C(A)$. Suppose g is a map from $A_1 * \dots * A_l$ to a torsion-free hyperbolic group G such that*

- *g is injective on edge groups of the graphs Λ_i ,*
- *if two edge groups of some of the graphs Λ_i have disjoint conjugacy classes in A_C , their images by g have disjoint conjugacy classes in G ;*
- *the images by g of intact surface type vertex groups are non-abelian.*

*Then there exists a finite union U_g of infinite cyclic subgroups of G such that for any map $\tau : (S_1 * \dots * S_p) * (Z_1 * \dots * Z_q) \rightarrow G$, if for all k , and for all j such that S_j is not a projective plane, the images $\tau(S_j)$ and $\tau(Z_k)$ are not contained in U_g , then the map $(g * \tau) \circ \rho_C : A \rightarrow G$ sends surface type vertex groups of Λ on non-abelian images.*

Proof. Let Σ be a surface of Λ , denote by S the corresponding surface type vertex group, and by v_S the corresponding vertex of Λ .

If $\Gamma(\Sigma, C)$ is a tree of groups, all of whose vertex groups except one are fundamental groups of spheres and projective planes, then $\rho_C(S)$ contains an intact surface type vertex group Q_i . The image of Q_i by g is non-abelian, so the image of S by $(g * \tau) \circ \rho_C$ is non-abelian regardless of the choice of τ . We may now assume that $\Gamma(\Sigma, C)$ is not a tree all of whose vertex groups except one are trivial or $\mathbb{Z}/2\mathbb{Z}$.

Suppose that S has two maximal boundary subgroups B_1 and B_2 whose conjugacy classes are disjoint in A . The edges adjacent to v_S corresponding to B_1 and B_2 are not adjacent to a same Z type vertex. Thus the groups $\rho_C(B_1)$ and $\rho_C(B_2)$ stabilise two edges which are not adjacent to a same Z type vertex. If these two edges lie in the same Λ_i , by remark 7.10 they have edge groups whose conjugacy classes are disjoint in $\rho_C(A)$ since Λ_i is a JSJ-like decomposition. If they lie in distinct graphs Λ_i , the corresponding edge groups lie in conjugates of distinct free factors of $\rho_C(A)$, so they also have disjoint conjugacy classes. By our second assumption on g , regardless on the choice of the map τ , the image of S by $(g * \tau) \circ \rho_C$ contains two cyclic subgroups whose conjugacy classes are disjoint. In particular it is not cyclic, and thus it is non-abelian since it lies in G which is torsion-free hyperbolic.

We may thus assume by remark 7.10 that all the edges adjacent to v_S in Λ are adjacent to a common Z -type vertex. This implies that all the edges in Λ_C adjacent to the subgraph $\Gamma(\Sigma, C)$ are adjacent to a same vertex w_Z , whose group is infinite cyclic. Let z be a generator of the group corresponding to w_Z .

Suppose first that Σ has at least two boundary components, with corresponding boundary subgroups B_1 and B_2 . Recall that in the graph of groups $\Delta(\Sigma, C)$, the groups B_1 and B_2 are elliptic, thus their images $\rho_C(B_1)$ and $\rho_C(B_2)$ stabilise vertices w_1 and w_2 of $\Gamma(\Sigma, C)$. There is a (possibly trivial) path in $\Gamma(\Sigma, C)$ joining the two vertices w_1 and w_2 . This path, together with the two edges joining w_1 and w_2 to w_Z gives a loop in the graph of groups Λ_C , that can be chosen to contain exactly one edge which is not in the maximal subtree we chose to define our pinching decomposition. Thus, up to replacing S by a conjugate, $\rho_C(S)$ contains z and tzt^{-1} , where t is a generator of one of the factors Z_j of the pinching decomposition of ρ_C . If $\tau(t)$ does not lie in the maximal cyclic subgroup $C_{g(z)}$ containing $g(z)$, then $(g * \tau)(z)$ and $(g * \tau)(tzt^{-1})$ do not commute in G , so $(g * \tau) \circ \rho_C(S)$ is not abelian.

Suppose now Σ has only one boundary component. Then either $\Gamma(\Sigma, C)$ is not a tree of groups, or it contains an interior vertex whose group is the fundamental group of a closed surface of positive genus. Up to replacing S by a conjugate, we see that $\rho_C(S)$ contains both z and either

one of the factors Z_j , or one of the factors S_j of the pinching decomposition. If the image of this factor by τ lies outside of the maximal cyclic subgroup $C_{g(z)}$ containing $g(z)$, the image of S by $(g * \tau) \circ \rho_C$ is not abelian. \square

Note that we have

Lemma 10.10: *If G is a torsion-free hyperbolic group, a non-cyclic subgroup G' of G is not contained in a finite union of cyclic groups.*

Proof. Cyclic subgroups of G are quasiconvex, so such a reunion U has growth in G at most linear (that is, the size of the set $B_G(n) \cap U$ grows linearly with n). Since G' is torsion-free and non-abelian, its growth is non-linear, thus the size of the set $B_G(n) \cap G'$ grows faster than any linear function. \square

Remark 10.11: *If G is torsion-free hyperbolic, if $g : A_C \rightarrow G$ sends intact surface type vertex groups on non-abelian images, is injective on edge groups of Λ_i , and preserves disjointness of conjugacy classes of edge groups, and if $g(A_1 * \dots * A_l)$ is not cyclic, we can always find a map $\tau : (S_1 * \dots * S_p) * (Z_1 * \dots * Z_q) \rightarrow g(A_1 * \dots * A_l)$ such that the map $(g * \tau) \circ \rho_C : A_C \rightarrow g(A_1 * \dots * A_l)$ sends surface type vertex groups on non-abelian images.*

The following lemma shows in particular that if a preretraction $f : A \rightarrow G$ factors as $f' \circ \rho_C$ where C is a maximal set of essential curves killed by f , and G is torsion-free hyperbolic, then intact surface type vertex groups have non-abelian images by f' , so f' satisfies the conditions of lemma 10.9.

Lemma 10.12: *Let $f : A \rightarrow G$ be a morphism which sends surface type vertex groups of Λ onto non-abelian images, and is injective on edge groups. Let C be a maximal set of essential curves killed by f , so that f factors as $f' \circ \rho_C$. Suppose G is torsion-free hyperbolic.*

If S is a surface type vertex group corresponding to an intact surface Σ of Λ_i , and if $\Delta(\Sigma, C^+)$ is a graph of group decomposition dual to a set of essential curves C^+ on Σ , there is at least one vertex group of $\Delta(\Sigma, C^+)$ whose image by f' is non-abelian.

Proof. We show first that $f'(S)$ is non-abelian. The group S is the unique infinite vertex group of one of the graph of groups of the form $\Gamma(\Sigma_0, C)$ for some surface Σ_0 of Λ , and we know that the graph underlying $\Gamma(\Sigma_0, C)$ is a tree of groups. Since G is torsion-free, the image by f' of the other finite vertex groups of $\Gamma(\Sigma_0, C)$ are trivial, so that the image of the fundamental group $\rho_C(S_0)$ of $\Gamma(\Sigma, C)$ by f' is exactly the image of S by f' : we have $f(S_0) = f'(\rho_C(S_0)) = f'(S)$. Now since S_0 is a surface type vertex group of Λ , its image by f is non-abelian, which proves the claim.

Suppose now all the vertex groups of $\Delta(\Sigma, C^+)$ have abelian image by f' (thus infinite cyclic since G is hyperbolic). Since f' is non-pinching with respect to Λ_i , the edge groups of $\Delta(\Sigma, C^+)$ are sent injectively into G by f' . This gives a graph of group decomposition of $f'(S)$ all of whose vertex and edge groups are infinite cyclic, so $f'(S)$ is a generalised Baumslag-Solitar group. In a generalised Baumslag-Solitar group, the commensurator of an elliptic element is the whole group (see for example [For02]). But in a torsion-free hyperbolic group, commensurators of elements are cyclic groups. This contradicts the non-abelianity of $f'(S)$, thus at least one of the vertex groups of $\Delta(\Sigma, C^+)$ has non-abelian image by f' . \square

10.3 Maximal preretractions

For the rest of this section, we let A be a torsion-free hyperbolic group which admits a cyclic JSJ-like decomposition Λ , and we assume that there exists at least one non-injective preretraction $A \rightarrow A$ with respect to Λ .

Definition 10.13: (set $L(f)$) If $f : A \rightarrow A$ is a preretraction, we denote by $L(f)$ the set of surfaces of Λ such that for at least one of the corresponding vertex groups S , the intersection $f(A) \cap S$ is not contained in a boundary subgroup of S .

Consider the set of tuples $(f, \mathcal{C}, \mathcal{C}^+)$ for which

- f is a non-injective preretraction $A \rightarrow A$;
- \mathcal{C} is a maximal set of curves on the surfaces of Λ killed by f , so that there exists $f' : \rho_{\mathcal{C}}(A) \rightarrow A$ with $f = f' \circ \rho_{\mathcal{C}}$;
- \mathcal{C}^+ is a set of essential curves on the surfaces of the graph of groups Λ_i obtained in the pinching of Λ by \mathcal{C} , such that \mathcal{C}^+ gives elliptic refinements Λ_i^+ of each Λ_i relatively to f' and Λ .

We say that an element $(f, \mathcal{C}, \mathcal{C}^+)$ is greater than another element $(g, \mathcal{D}, \mathcal{D}^+)$ if \mathcal{C} strictly contains \mathcal{D} , or if they are equal and \mathcal{C}^+ strictly contains \mathcal{D}^+ , or if they too are equal, and $L(f)$ is contained in $L(g)$ (note the inversion).

A preretraction f for which there exists \mathcal{C} and \mathcal{C}^+ such that $(f, \mathcal{C}, \mathcal{C}^+)$ is a maximal element in our set is called a maximal non-injective preretraction. Such an element must exist, indeed, the set of non-injective preretractions is not empty, the cardinal of a set of essential curves on a finite set of surfaces is bounded, and the set $L(f)$ is a subset of the finite set of surfaces of Λ .

For the rest of this section, we let $f : A \rightarrow A$ be a maximal non-injective preretraction for the sets of curves \mathcal{C} and \mathcal{C}^+ . Build the pinching of Λ by \mathcal{C} , a pinching decomposition of $\rho_{\mathcal{C}}(A)$, and elliptic refinements Λ_i^+ of the graphs of groups Λ_i given by \mathcal{C}^+ . By remark 9.12, we have minimal equivariant maps $t_i^+ : T_{\Lambda_i^+} \rightarrow T_{\Lambda}$.

A very important property of such a maximal element is given by

Lemma 10.14: For any surface Σ of Λ , the following are equivalent:

- (i) $\Sigma \in L(f)$;
- (ii) one of the surface type vertex v of T_{Λ} corresponding to Σ lies in the image of one of the maps $t_i^+ : T_{\Lambda_i^+} \rightarrow T_{\Lambda}$;
- (iii) for one of the surface type vertex group S corresponding to Σ , there is a surface type vertex group S^+ of one of the elliptic refinements Λ_i^+ such that $f'(S^+)$ is a subgroup of finite index S .

Proof. The equivalence between (ii) and (iii) is given by lemma 9.14. It is clear that (iii) implies (i). Let us show that the converse is true.

If $m = 1$, there is only one component Λ_1 , the result is given by lemma 9.14.

If $m \geq 2$, \mathcal{C} is not empty, and by lemma 10.4, the surface complexity of each of the graph of groups Λ_i is smaller than the surface complexity of Λ . Consider the map $h = f'|_{A_{\mathcal{C}}} : A_1 * \dots * A_m \rightarrow A$. The hypotheses of lemma 9.18 are satisfied, so we can find a map $\tilde{h} : A_1 * \dots * A_m \rightarrow A$ such that $\tilde{h}|_{A_i}$ coincides with $f'|_{A_i}$ up to conjugation, and if S is a surface type vertex group of Λ whose intersection with $\tilde{h}(A_1 * \dots * A_m)$ is not contained in a boundary subgroup of S , then a conjugate of S contains with finite index the image of a surface type vertex group S^+ of one of the graphs of groups Λ_i^+ . We also know that $\tilde{h}(A_1 * \dots * A_m)$ is the free product of the $\tilde{h}(A_i)$, so in particular it is not abelian since we assumed $m \geq 2$.

By lemma 10.12, the map f' sends intact surface type vertex groups of the graph of groups Λ_i to non-abelian images. The image by \tilde{h} of an intact surface type vertex group S of one of the graphs Λ_i is just a conjugate of $f'(S)$, so it is also non-abelian. Thus by remark 10.11, there exists a map $\tau : R_{\mathcal{C}} \rightarrow \tilde{h}(A_1 * \dots * A_m)$ such that the map $F = (\tilde{h} * \tau) \circ \rho_{\mathcal{C}}$ sends surface type vertex

groups of Λ to non-abelian images. We now want to see that $(F, \mathcal{C}, \mathcal{C}^+)$ is a maximal non-injective preretraction.

It is easy to check that F restricts to a conjugation on each non surface type vertex group of Λ , so that F is a preretraction. The map F factors through $\rho_{\mathcal{C}}$, so F is in fact a non-injective preretraction, and the curves of \mathcal{C} are killed by F . By maximality of f , we see that \mathcal{C} is a maximal set of curves killed by F . The map f' sends elements corresponding to curves of \mathcal{C}^+ to edge groups of Λ , thus so does the map \tilde{h} . Similarly, by maximality of f , the curves \mathcal{C}^+ must give elliptic refinements of the graph of groups Λ_i with respect to $(\tilde{h} * \tau)$ and Λ .

Now, let S be a surface type vertex group of Λ whose corresponding surface is in $L(F)$. Its intersection with $F(A) = \tilde{h}(A_1 * \dots * A_l)$ is not contained in a boundary subgroup, so by our choice of \tilde{h} , there is a surface type vertex group S^+ of one of the graphs of groups Λ_i^+ such that $\tilde{h}(S^+)$ is a subgroup of finite index of S . But on A_i , the maps \tilde{h} and f' coincide up to conjugation: thus $f'(S^+)$ is a subgroup of finite index of some conjugate of S . We have shown that to any surface Σ which lies in $L(F)$ corresponds a group S which admits as a subgroup of finite index the image by f' of a surface type vertex group S^+ of Λ_i^+ .

This implies first that $L(F) \subseteq L(f)$. By maximality of f , we see that this must in fact be an equality. But then if Σ is in $L(f)$, it is also in $L(F)$, so there is a group S with corresponding surface Σ which admits as a subgroup of finite index the image by f' of a surface type vertex group S^+ of Λ_i^+ : we see that (iii) must hold. \square

From this we deduce in particular

Lemma 10.15: *The set $L(f)$ does not contain all the surfaces of Λ .*

Proof. Suppose that $L(f)$ contains all the surfaces of Λ . By lemma 10.14, for every surface Σ of Λ , there exists a surface type vertex group S^+ of one of the graphs of groups Λ_i^+ such that $f(S^+)$ is a subgroup of finite index of one a surface type vertex group S corresponding to Σ . Moreover, f' sends boundary subgroups of S^+ to boundary subgroups of S . By lemma 9.10, the complexity of the surface Σ^+ corresponding to S^+ is greater than or equal to that of Σ , and if we have equality, $f'|_{S^+}$ is an isomorphism onto S . This implies that the complexity of the set of all the surfaces of the Λ_i^+ is greater than the surface complexity of Λ . But lemma 9.16 shows that the set of all the surfaces of the Λ_i^+ has complexity smaller than the set of surfaces of the Λ_i , which in turn has complexity smaller than the surface complexity of Λ by lemma 10.4. Thus these complexities are all equal, which implies that the sets \mathcal{C} and \mathcal{C}^+ are empty, and that each surface type vertex group of Λ is sent isomorphically onto a surface type vertex group of Λ by f , non-conjugate surface type vertex groups being sent to non-conjugate surface type vertex group.

Thus some power of f sends each surface type vertex group of Λ isomorphically on a conjugate of itself, and restricts to conjugation on each non surface type vertex group. By proposition 8.1, it is an isomorphism. This contradicts the non-injectivity of f . \square

We now want to define applications $P(f, k) : A \rightarrow A$ which we call pseudo-powers of f . Indeed, we need to iterate f , but we want the result to still be a preretraction, this is why we cannot take simply the powers of f since they might send surface type vertex groups onto abelian images.

We define $P(f, k)$ by induction as follows. Let $P(f, 1) = f$. If $P(f, k-1)$ is defined, and is a maximal preretraction $A \rightarrow A$ we consider the map $P(f, k-1) \circ (f'|_{A_{\mathcal{C}}}) : A_{\mathcal{C}} \rightarrow A$.

Lemma 10.16: *The map $P(f, k-1) \circ (f'|_{A_{\mathcal{C}}})$ sends intact surface type vertex groups of the graphs of groups Λ_i to non-abelian images.*

Proof. If S is an intact surface type vertex of one of the graphs of groups Λ_i with corresponding surface Σ , it inherits a decomposition $\Delta(\Sigma, \mathcal{C}^+)$ from the elliptic refinement Λ_i^+ . We know by lemma 10.12 that there is at least one of the vertex groups S_0 of $\Delta(\Sigma, \mathcal{C}^+)$ whose image by f' is non-abelian. If $f'(S_0)$ lies in a non-surface type vertex group of Λ , the preretraction $P(f, k-1)$

is injective on $f'(S_0)$, so $P(f, k-1) \circ f'(S)$ is non-abelian. If $f'(S_0)$ lies in a surface type vertex group S_1 of Λ , it must be with finite index by lemma 9.7 since f' is non-pinching on Σ and sends boundary elements on edge groups of Λ . Now $P(f, k-1) \circ f'(S_0)$ is a subgroup of finite index of $P(f, k-1)(S_1)$, which is non-abelian since $P(f, k-1)$ is a preretraction. Thus $P(f, k-1) \circ f'(S)$ is non-abelian. \square

We will now build $P(f, k)$. Since A is non-abelian and torsion-free hyperbolic, Λ admits at least one non-abelian vertex group: if it is a non surface type vertex group, $P(f, k-1)$ sends it injectively into A so its image is non-abelian, and if it is a surface vertex group its image by $P(f, k-1)$ is non-abelian by definition of a preretraction. This shows that $P(f, k-1)(A)$ is not cyclic. We can thus apply remark 10.11 to $P(f, k-1) \circ f'|_{A_C}$, this tells us we can find a map $\tau : R_C \rightarrow P(f, k-1)(A)$ such that the map

$$P(f, k) = [(P(f, k-1) \circ f'|_{A_C}) * \tau] \circ \rho_C$$

sends surface type vertex groups on non-abelian images. Let us now see that $P(f, k)$ is a maximal non-injective preretraction. It is easy to see that $P(f, k)$ sends non surface type vertex groups on conjugates of themselves, so is in fact a preretraction. If \mathcal{C} is empty, $f' = f$ so $P(f, k)$ is not injective since f is not injective.

Since $P(f, k)$ factors through ρ_C , it kills the curves in \mathcal{C} , so by maximality of $P(f, k-1)$ the set \mathcal{C} is a maximal set of essential curves killed by $P(f, k)$. Similarly since $P(f, k-1)$ conjugates edge groups, $P(f, k-1) \circ f'|_{A_C}$ sends elements corresponding to curves of \mathcal{C}^+ to edge groups of Λ , so by maximality of $P(f, k-1)$, the set \mathcal{C}^+ is a maximal set of essential curves that give an elliptic refinement of the graphs of groups Λ_i with respect to $P(f, k-1) \circ f'|_{A_C}$ and Λ . Finally, the image of $P(f, k)$ is contained in the image of $P(f, k-1)$, so $L(P(f, k)) \subseteq L(P(f, k-1))$, and by maximality of $P(f, k-1)$ this is in fact an equality. Thus $P(f, k)$ is a maximal non-injective preretraction.

Using pseudo-powers, we can now show

Lemma 10.17: *If f is a non-injective preretraction, it sends each surface type vertex group corresponding to a surface of $L(f)$ isomorphically onto another surface group corresponding to a surface of $L(f)$.*

Proof. We have just seen that $P(f, 2)$ is also a maximal preretraction, for the same sets \mathcal{C} and \mathcal{C}^+ . Thus $P(f, 2)$ factors through ρ_C , we write $P(f, 2) = [P(f, 2)]' \circ \rho_C$. Recall that $P(f, 2) = [(f \circ f') * \tau] \circ \rho_C$ so that $[P(f, 2)]'|_{A_C} = f \circ f'$. Let Σ be a surface of $L(f)$.

Since $L(f) = L(P(f, 2))$, and using lemma 10.14, we see that there is a group S with corresponding surface Σ , and a surface type vertex group S^+ of Λ_i^+ for some i , such that $f \circ f'(S^+)$ is a subgroup of finite index of S . Consider $f'(S^+)$: it is elliptic in Λ since the Λ_i^+ are elliptic refinements relative to both f' and $[P(f, 2)]'$.

It cannot lie in a non surface type vertex group of Λ , since these are sent to conjugates of themselves by f . Thus it lies in a surface type vertex group S_1 of Λ , and by lemma 9.7, it is a subgroup of finite index of S_1 . This implies in particular that the surface corresponding to S_1 is in $L(f)$.

Now $f(S_1)$ contains a subgroup of finite index, namely $f(f'(S^+))$, which is elliptic in T_Λ : thus $f(S_1)$ itself is elliptic. Thus it lies in a vertex stabiliser of T_Λ , which must in fact be S . Since $f(f'(S^+))$ has finite index in S , so does $f(S_1)$. By lemma 9.10, the complexity of the surface Σ_1 corresponding to S_1 is greater than that of the surface Σ corresponding to S , and if we have equality, $f|_{S_1}$ is an isomorphism onto S .

Thus to each surface Σ in $L(f)$ corresponds a surface Σ_1 in $L(f)$ whose complexity is greater, and such that any group S_1 corresponding to Σ_1 has image by f lying in a group S corresponding to Σ . In particular, the map $\Sigma \mapsto \Sigma_1$ is injective. Since it is a map $L(f) \rightarrow L(f)$, it is a bijection,

thus for any surface Σ of $L(f)$, we must have $k(\Sigma) = k(\Sigma_1)$. This implies that f sends each surface type vertex group whose surface is in $L(f)$ isomorphically onto a surface type vertex group whose surface is in $L(f)$. \square

Remark 10.18: *The previous lemma implies that some pseudo-power $P(f, k)$ of the map f sends surface type vertex groups whose corresponding surface is in $L(f)$ isomorphically onto conjugates of themselves. Thus, there exists a maximal non-injective preretraction f which sends surface type vertex groups corresponding to surfaces of $L(f)$ isomorphically onto conjugates of themselves.*

10.4 Proof of Proposition 7.15

Again, A is a torsion-free hyperbolic group which admits a cyclic JSJ-like decomposition Λ , and we assume that there exists at least one non-injective preretraction $A \rightarrow A$ with respect to Λ . We let f be a maximal non-injective preretraction which sends surface type vertex groups whose corresponding surface are in $L(f)$ isomorphically onto conjugates of themselves, as is given by remark 10.18.

Consider the complement in Λ of the set containing surface type vertices corresponding to surfaces which do *not* lie in $L(f)$, as well as the open edges adjacent to these vertices. Its connected components $\Gamma_1, \dots, \Gamma_m$ are subgraphs of groups of Λ , we denote their fundamental groups by H_1, \dots, H_m .

Call Γ the graph of groups with surfaces obtained by collapsing in Λ all the edges of the subgraphs Γ_i . If we choose a maximal subtree in Γ , as well as a lift to the corresponding tree T_Γ , we identify the groups H_i to subgroups of A . Given a preferred non-surface type vertex R_0 , we can do this in such a way that R_0 lies in one of the subgroups H_i . The subgroup of A generated by H_1, \dots, H_m will be our retract A' .

Note that T_Γ is bipartite, in the sense that any edge has one end whose vertex group is a conjugate of one of the subgroups H_i , and one end whose stabiliser is a surface type vertex group of Λ whose corresponding surface is not in $L(f)$.

Lemma 10.19: *The map f sends each H_i isomorphically onto a conjugate of itself.*

Proof. A vertex group of Γ_i is either a non-surface type vertex group, or a surface type vertex group whose conjugacy class is in $L(f)$: in both cases, it is sent isomorphically on a conjugate of itself in A by f . Now any two adjacent vertex groups G_v and G_w of Γ_i intersect in a non-trivial edge group, and since f is injective on edge groups, the intersection $f(G_v) \cap f(G_w)$ contains a non-trivial element. If $f(G_v) = g_v G_v g_v^{-1}$ and $f(G_w) = g_w G_w g_w^{-1}$, the intersection $g_v H_i g_v^{-1} \cap g_w H_i g_w^{-1}$ contains $f(G_v) \cap f(G_w)$, so in particular it is non-trivial: again by bipartism of Γ and 1-acylindricity near surface type vertices, we deduce that $g_v H_i g_v^{-1} = g_w H_i g_w^{-1}$ so that $g_w^{-1} g_v$ is in H_i , and $f(G_w)$ and $f(G_v)$ lie in the same conjugate of H_i . Thus, $f|_{H_i}$ composed by the conjugation by g_v^{-1} restricts to a conjugation by an element of H_i on non-surface type vertex groups of Γ_i , and sends surface type vertex groups isomorphically on conjugates of themselves by an element of H_i .

As fundamental groups of subgraphs of groups of Λ , the groups H_i are quasiconvex in A , thus they are hyperbolic. Note also that the decomposition Γ_i is a JSJ-like decomposition for H_i . We can now apply proposition 8.1 to conclude that $f|_{H_i}$ composed with the conjugation by g_v^{-1} is an isomorphism $H_i \rightarrow H_i$. Thus f itself sends each H_i isomorphically onto a conjugate of itself: the claim is proved. \square

Recall we chose a pinching decomposition of $\rho_C(A)$, and we let f' be such that $f = f' \circ \rho_C$. Recall also that the set \mathcal{C}^+ gave us elliptic refinements Λ_i^+ for each Λ_i with respect to f' and Λ .

Lemma 10.20: *For each j , the image $f'(A_j)$ lies in a conjugate of one of the subgroups H_i .*

Proof. For each Λ_j^+ , we have a minimal equivariant map $t_j^+ : T_{\Lambda_j^+} \rightarrow T_\Lambda$. By lemma 10.14, the image of t_j^+ contains none of the vertices corresponding to surfaces which are not in $L(f)$. Since $t_j^+(T_{\Lambda_j^+})$ is connected, this implies that the image of A_j by f lies in a conjugate of one of the subgroups H_i . \square

Fix an index i . It is straightforward to see that $\rho_C(H_i)$ lies in a conjugate of one of the subgroups A_{j_i} . We just saw that H_i is sent isomorphically onto a conjugate of itself by f , thus A_{j_i} must be sent to a conjugate of H_i by f' . In particular, the application $i \mapsto j_i$ is injective. Conversely, each A_j contains a conjugate of one of the $\rho_C(H_i)$. Up to renumbering, we may thus assume that $\rho_C(H_i)$ is contained in a conjugate of A_i .

Lemma 10.21: *The group A' generated by H_1, \dots, H_m is the free product $H_1 * \dots * H_m$.*

Proof. Recall that the group A_C generated by the groups A_i is in fact the free product of the groups A_i . Since the A_i form a free product in $\rho_C(A)$, the group $\rho_C(A')$ generated by the subgroups ρ_{H_i} is in fact the free product of the subgroups $\rho_C(H_i)$. Since ρ_C is injective on H_i , this means that the H_i themselves form a free product. \square

Note that since the list $L(f)$ does not contain all the surfaces of Λ , the group A' is a proper subgroup of A .

We now want to understand the image of $f'|_{A_C}$. For each i , we have $f'(A_i) = g_i H_i g_i^{-1}$. The image of A_C by f' is generated by these conjugates of the subgroups H_i . It acts on the tree T_Γ corresponding to Γ . A surface type vertex group S of Γ corresponds to a surface which does not lie in $L(f)$, so it intersects $f'(A_C)$ at most in a boundary subgroup. Thus, in the action of $f'(A_C)$, the corresponding vertex has cyclic stabiliser, and if it is not trivial, it stabilises an adjacent edge. This edge is unique by 1-acylindricity of surface type vertices, so by collapsing all such edges, we see that

$$f'(A_C) = g_1 H_1 g_1^{-1} * g_2 H_2 g_2^{-1} * \dots * g_l H_l g_l^{-1}.$$

Let β be the map which restricts on $g_i H_i g_i^{-1}$ to conjugation by g_i^{-1} . The map f' sends intact surface type vertex groups to non-abelian images by remark 10.12, hence so does $\beta \circ f'$ since β is an isomorphism between $g_1 H_1 g_1^{-1} * \dots * g_l H_l g_l^{-1}$ and $H_1 * \dots * H_l$.

If $l \geq 1$, then $H_1 * \dots * H_l$ is clearly non-cyclic. But if $l = 1$, the image of A by f is contained in H_1 , so H_1 is not abelian. Thus, by lemma 10.11, we can find a map $\tau : R_C \rightarrow \beta(f'(A_C)) = H_1 * \dots * H_l$ such that the map $F = [(\beta \circ f') * \tau] \circ \rho_C$ sends surface type vertex groups to non-abelian images. Moreover, it is easy to see that the map F sends each subgroup H_i isomorphically on itself.

Thus the restriction η of F to $A' = H_1 * \dots * H_m$ is an isomorphism $A' \rightarrow A'$. Finally, the map $\eta^{-1} \circ F$ is a retraction r from A to A' , which sends surface type vertex groups of Γ to non-abelian images. Now A admits a graph of groups decomposition with one non-surface type vertex v stabilised by A' , the other vertices being stabilised by surface type vertex groups of Γ , and edges joining these to v . Thus (A, A', r) is a hyperbolic floor. This terminates the proof of proposition 7.15.

10.5 Proof of Proposition 7.16

Let A be a group which admits a JSJ-like decomposition Λ . Suppose G' is a subgroup of G containing A such that either G' is a free factor of G , or G' is a retract of G by a retraction $r : G \rightarrow G'$ which makes (G, G', r) a hyperbolic floor.

Denote by $r : G \rightarrow G'$ the retraction which is the trivial map on R if $G = G' * R$, and the retraction of the hyperbolic floor structure in the second case. Let Γ be the graph of group

corresponding to the free product $G' * R$ in the first case, and the graph of groups decomposition associated to the hyperbolic floor structure in the second case.

Let $f : A \rightarrow G$ be the preretraction given by the hypotheses. Choose a maximal system \mathcal{C} of essential curves killed by f on the surfaces of Λ . Let f' be such that $f = f' \circ \rho_{\mathcal{C}}$. Build the pinching of Λ by \mathcal{C} , and choose a pinching decomposition of $\rho_{\mathcal{C}}(A)$. Choose also a maximal system of essential curves \mathcal{C}^+ which gives an elliptic refinement Λ_i^+ for each Λ_i , with respect to f' and Γ .

Lemma 10.22: *The map $r \circ f'|_{A_{\mathcal{C}}}$ sends intact surface type vertex groups of the Λ_i to non-abelian images.*

Proof. Let S be an intact surface type vertex group of Λ_i . It inherits a decomposition $\Delta(\Sigma, \mathcal{C}^+)$ from the refinement Λ_i^+ .

If we are in the case where $G = G' * R$, the set of curves of \mathcal{C}^+ lying on the surface Σ corresponding to S is empty. Indeed, elements corresponding to curves of \mathcal{C}^+ are sent to edge groups of Γ by f' , but edge groups of Γ are trivial and f' is non-pinching with respect to Λ_i so there can be no curves of \mathcal{C}^+ on Σ . Thus $f'(S)$ is elliptic in Γ . Since boundary subgroups of S are sent to non-trivial subgroups of a conjugate of A , $f'(S)$ lies in a conjugate of G' , and by lemma 10.12, it is non-abelian. Thus its image by r is non-abelian.

Let us now assume we are in the case where $(G, r(G), r)$ is a hyperbolic floor. By lemma 10.12, the image of at least one of the vertex groups S^+ of $\Delta(\Sigma, \mathcal{C}^+)$ has non-abelian image by f' . If $f'(S^+)$ lies in a conjugate of G' , its image by r is clearly non-abelian. If $f'(S^+)$ lies in one of the surface type vertex groups S_1 of T_{Γ} , it is a subgroup of finite index of S_1 by lemma 9.7. Now this means $r(f'(S^+))$ is a finite index subgroup of $r(S_1)$, which is not abelian by definition of a hyperbolic floor. Hence it is itself non-abelian. \square

Note that G' contains A , so that it isn't cyclic. Now we can apply remark 10.11 to $r \circ f'|_{A_{\mathcal{C}}}$, to get a map $\tau : R_{\mathcal{C}} \rightarrow G'$ such that the map $[(r \circ f'|_{\mathcal{C}}) * \tau] \circ \rho_{\mathcal{C}}$ sends surface type vertex groups of Λ on non-abelian images. It is easy to see that this map restricts to conjugation on each non surface type vertex group of Λ . This shows precisely that it is a preretraction $A \rightarrow G'$. If \mathcal{C} is not empty, $\rho_{\mathcal{C}}$ is not injective, thus so is $[(r \circ f'|_{\mathcal{C}}) * \tau] \circ \rho_{\mathcal{C}}$. If \mathcal{C} is empty, $[(r \circ f'|_{\mathcal{C}}) * \tau] \circ \rho_{\mathcal{C}}$ is just $r \circ f$ so it is also non-injective. This terminates the proof of proposition 7.16.

Bibliography

- [BF] Mladen Bestvina and Mark Feighn, *Outer Limits*, preprint, available at <http://andromeda.rutgers.edu/~feighn/papers/outer.pdf>.
- [BF95] ———, *Stable actions of groups on real trees*, Invent. Math. **121** (1995), 287–321.
- [BF03] ———, *Notes on Sela’s work: Limit groups and Makanin-Razborov diagrams*, Geometry and Cohomology in Group Theory (2003).
- [CDP90] Michel Coornaert, Thomas Delzant, and Athanase Papadopoulos, *Géométrie et théorie des groupes: les groupes hyperboliques de Gromov*, vol. 1441, Springer-Verlag, 1990.
- [CG05] Christophe Champetier and Vincent Guirardel, *Limit groups as limits of free groups: compactifying the set of free groups*, Israel J. Math. **146** (2005), 1–76.
- [Cha] Zoé Chatzidakis, *Introduction to model theory*, Notes, available at <http://www.logique.jussieu.fr/~zoe/index.html>.
- [Del96] Thomas Delzant, *Quotients des groupes hyperboliques*, Duke Mathematical Journal **83** (1996), 661–682(3).
- [DS99] Martin Dunwoody and Michah Sageev, *JSJ-splittings for finitely presented groups over slender groups*, Invent. Math. **135** (1999), 25–44.
- [FLP79] A Fathi, F Laudenbach, and Valentin Poenaru, *Travaux de Thurston sur les surfaces*, vol. 66-67, Soc. Math. Fr., 1979.
- [For02] Max Forester, *Deformation and rigidity of simplicial group actions on trees*, Geometry and Topology **6** (2002), 219–267.
- [FP06] Koji Fujiwara and Panos Papasoglu, *JSJ-decompositions of finitely presented groups and complexes of groups*, GAFA **16** (2006), 70–125.
- [GdlH90] Étienne Ghys and Pierre de la Harpe, *Sur les groupes hyperboliques d’après Mikhael Gromov*, Birkhäuser, 1990.
- [Gro05] Daniel Groves, *Limit groups for relatively hyperbolic groups, II: Makanin-Razborov diagrams*, Geom. Top. **9** (2005), 2319–2358.
- [Gui08] Vincent Guirardel, *Action of finitely generated groups on \mathbb{R} -trees*, Annales de l’Institut Fourier **58** (2008), 159–211.
- [KM98] Olga Kharlampovich and Alexei Myasnikov, *Irreducible affine varieties over a free group II*, J. Algebra **200** (1998), 517–570.
- [KM06] ———, *Elementary theory of free nonabelian groups*, J. Algebra **302** (2006), 451–552.

- [Lev94] Gilbert Levitt, *Graphs of actions on \mathbb{R} -trees*, Comment. Math. Helv. **69**(1) (1994), 28–38.
- [Lic64] W.B.R. Lickorish, *A finite set of generators for the homeotopy group of a 2-manifold*, Proc. Cambridge Philos. Soc. **60** (1964), 769–778.
- [LP97] Gilbert Levitt and Frédéric Paulin, *Geometric group actions on trees*, Amer. J. Math. **119**(1) (1997), 83–102.
- [MS84] Morgan and Peter Shalen, *Degeneration of hyperbolic structures*, Ann. of Math. **120** (1984), 401–476.
- [MS91] ———, *Free actions of surface groups on R -trees*, Topology **30** (1991), 143–154.
- [Pau88] Frédéric Paulin, *Topologie de Gromov équivariante, structures hyperboliques et arbres réels*, Invent. Math. **94** (1988), 53–80.
- [Pau89] ———, *The Gromov topology on R -trees*, Topology Appl. **32** (1989), 197–221.
- [Pau97] ———, *On outer automorphisms of hyperbolic groups*, Annales Scientifiques de l'Ecole Normale Supérieure **30** (1997), 147–167(21).
- [Pau03] ———, *Sur la théorie élémentaire des groupes libres*, Séminaire Bourbaki **922** (2002–2003), 1–39.
- [Raz85] Razborov, *On systems of equations in a free group*, Mathematics of the USSR-Izvestiya **25** (1985), 115–162.
- [RS94] Eliyahu Rips and Zlil Sela, *Structure and rigidity in hyperbolic groups I*, Geom. Funct. Anal. **4** (1994), 337–371.
- [RS97] ———, *Cyclic splittings of finitely presented groups and the JSJ decomposition*, Ann. of Math. **146** (1997), 53–104.
- [Sco78] Peter Scott, *Subgroups of surface groups are almost geometric*, Journal of the London Mathematical Society **17** (1978), 555–565.
- [Sel] Zlil Sela, *Diophantine geometry over groups VII: The elementary theory of a hyperbolic group.*, Proc. London Math Soc., to appear.
- [Sel97a] ———, *Acyindrical accessibility*, Invent. Math. **129** (1997), 527–565.
- [Sel97b] ———, *Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups II*, Geom. Funct. Anal. **7** (1997), 561–593.
- [Sel01] ———, *Diophantine geometry over groups I: Makanin-Razborov diagrams*, Publications Mathématiques de l'IHÉS **93** (2001), 31–105.
- [Sel03] ———, *Diophantine geometry over groups II: Completions, closures and formal solutions.*, Israel J. Math. **134** (2003), 173–254.
- [Sel04] ———, *Diophantine geometry over groups IV: An iterative procedure for the validation of sentence.*, Israel J. Math. **143** (2004), 1–130.
- [Sel05a] ———, *Diophantine geometry over groups III: Rigid and solid solutions.*, Israel J. Math. **147** (2005), 1–73.

- [Sel05b] ———, *Diophantine geometry over groups V: Quantifier elimination. Part I.*, Israel J. Math. **150** (2005), 1–198.
- [Sel06] ———, *Diophantine geometry over groups VI: The elementary theory of free groups.*, Geom. Funct. Anal. **16** (2006), 707–730.
- [Ser83] Jean-Pierre Serre, *Arbres, amalgames, SL_2* , Astérisque **46** (1983).
- [Wil] Henry Wilton, *Rips Theory*, Notes, available at <http://www.math.utexas.edu/users/henry.wilton/rips.pdf>.
- [Wil06] ———, *Subgroup separability of limit groups*, Ph.D. thesis, Imperial College, London, 2006, available at <http://www.math.utexas.edu/users/henry.wilton/thesis.pdf>.