

Rigidity of high dimensional graph manifolds

Roberto Frigerio

Jean-François Lafont

Alessandro Sisto

Author address:

DIPARTIMENTO DI MATEMATICA - LARGO B. PONTECORVO, 5 - 56127 PISA,
ITALY

E-mail address: `frigerio@dm.unipi.it`

DEPARTMENT OF MATHEMATICS, THE OHIO STATE UNIVERSITY, COLUM-
BUS, OH 43210, USA

E-mail address: `jlafont@math.ohio-state.edu`

MATHEMATICAL INSTITUTE, 24-29 ST GILES', OXFORD, OX1 3LB, ENG-
LAND, UK

E-mail address: `sisto@maths.ox.ac.uk`

Contents

Introduction	ix
Acknowledgments	xv
Part 1. Graph manifolds: topological and algebraic properties	1
Chapter 1. Quasi-isometries and quasi-actions	3
1.1. The quasi-isometry type of a group	3
1.2. The Milnor-Svarc Lemma	3
1.3. From quasi-isometries to quasi-actions	4
Chapter 2. Generalized graph manifolds	7
2.1. Putting a metric on a high dimensional graph manifold	8
2.2. $\pi_1(M)$ as the fundamental group of a graph of groups	9
2.3. The universal cover of M as a tree of spaces	10
2.4. Basic metric properties of \widetilde{M}	12
2.5. Examples not supporting any locally CAT(0) metric	13
Chapter 3. Topological rigidity	15
3.1. Contractible universal cover	15
3.2. Lower algebraic K-theory	16
3.3. Borel Conjecture	17
3.4. Baum-Connes Conjecture and consequences	20
Chapter 4. Isomorphisms preserve pieces	23
4.1. Edge stabilizers are preserved by isomorphisms	23
4.2. Isomorphisms quasi-preserve pieces	24
4.3. Isomorphisms preserve pieces	26
Chapter 5. Smooth rigidity	29
5.1. Defining ψ on the pieces	29
5.2. Dehn twists	31
5.3. Concluding the proof of smooth rigidity	33
5.4. Mapping class group	34
Chapter 6. Algebraic properties: general graph manifolds	37
6.1. Kazhdan subgroups	37
6.2. The Tits Alternative	38
6.3. Co-Hopf property	39
6.4. Uniformly exponential growth	42

Part 2. Irreducible graph manifolds: coarse geometric properties	45
Chapter 7. Irreducible graph manifolds	47
Chapter 8. Algebraic properties: irreducible graph manifolds	61
8.1. Acylindrical graph of groups	61
8.2. C^* -simplicity	62
8.3. SQ-universality	65
8.4. Solvable word problem	67
8.5. Gluings and isomorphism type	70
Chapter 9. Pieces of irreducible graph manifolds are quasi-preserved	77
9.1. The asymptotic cone of a geodesic metric space	77
9.2. Quasi-isometries and asymptotic cones	78
9.3. Tree-graded spaces	80
9.4. Studying the asymptotic cones of \widetilde{M}	80
9.5. Walls and chambers are quasi-preserved by quasi-isometries	84
9.6. Thickness and relative hyperbolicity	85
Chapter 10. Quasi isometry rigidity, I	87
10.1. The quasi-action of Γ on \widetilde{M}	87
10.2. The image of θ	89
10.3. The kernel of θ	91
10.4. Analyzing the exact sequence	93
10.5. Pieces with quasi-isometric fundamental groups	94
Chapter 11. Quasi isometry rigidity, II	97
11.1. From quasi-actions to actions on trees	97
11.2. The action of Γ^0 on T	98
11.3. Stabilizers of edges and vertices	98
11.4. Graph manifolds with quasi-isometric fundamental groups	99
Part 3. Concluding remarks	103
Chapter 12. Examples not supporting locally CAT(0) metrics	105
12.1. Fiber bundles	105
12.2. Irreducible examples	110
Chapter 13. Allowing surface pieces	115
13.1. Pieces with a surface base	115
13.2. Characterizing the wall subgroups	116
13.3. Rigidity results	122
Chapter 14. Directions for future research	125
14.1. Further algebraic properties	125
14.2. Studying quasi-isometries	126
14.3. Non-positive curvature and differential geometry	127
Bibliography	131

Abstract

We define the class of *high dimensional graph manifolds*. These are compact smooth manifolds supporting a decomposition into finitely many pieces, each of which is diffeomorphic to the product of a torus with a finite volume hyperbolic manifold with toric cusps. The various pieces are attached together via affine maps of the boundary tori. We require all the hyperbolic factors in the pieces to have dimension ≥ 3 (though several of our results can be extended to allow pieces with hyperbolic surface factors). Our main goal is to study this class of graph manifolds from the viewpoint of rigidity theory.

We show that, in high dimensions, the Borel conjecture holds for our graph manifolds. We also show that smooth rigidity holds within the class: two graph manifolds are homotopy equivalent if and only if they are diffeomorphic. We introduce the notion of *irreducible* graph manifolds. These form a subclass which has better coarse geometric properties, in that various subgroups can be shown to be quasi-isometrically embedded inside the fundamental group. We establish some structure theory for finitely generated groups which are quasi-isometric to the fundamental group of an irreducible graph manifold: any such group has a graph of groups splitting with strong constraints on the edge and vertex groups. Along the way, we classify groups which are quasi-isometric to the product of a free abelian group and a non-uniform lattice in $SO(n, 1)$ (including the $n = 2$ case, which corresponds to groups of the form $F_k \times \mathbb{Z}^d$). We provide various examples of graph manifolds which do *not* support any locally CAT(0) metric. Finally, we discuss how to adapt our arguments to the situation where some pieces are products of hyperbolic surfaces with codimension two tori.

We emphasize that, in dimension = 3, our notion of graph manifold **does not** coincide with the classical graph manifolds. Rather, it is a class of 3-manifolds that contains some (but not all) classical graph 3-manifolds (we don't allow general Seifert fibered pieces), as well as some non-graph 3-manifolds (we do allow hyperbolic pieces). Some of our results appear new even in the 3-dimensional case.

2000 *Mathematics Subject Classification*. Primary: 53C24, 20F65; Secondary: 53C23, 20E08, 20F67, 20F69, 19D35.

Key words and phrases. Quasi-isometry, quasi-action, graph of groups, CAT(0) space, Borel conjecture, Baum-Connes conjecture, asymptotic cone, mapping class group, Kazhdan's property (T), Tits' alternative, co-Hopf property, C^* -simplicity, SQ -universality.

Introduction

In recent years, there has been an extensive amount of work done on proving rigidity results for various classes of non-positively curved spaces. In this monograph, we are interested in establishing similar rigidity theorems in the context of spaces which may not support any non-positively curved metrics.

To motivate our class of manifolds, we briefly recall some basic notions from 3-manifold topology. In the theory of 3-manifolds, a central role is played by *Thurston's geometrization conjecture*, recently established by Perelman. Loosely speaking, this asserts that a closed 3-manifold can be decomposed into pieces, each of which supports a *geometric structure*, i.e. a complete metric locally modelled on one of the eight 3-dimensional geometries. When restricted to the class of 3-manifolds which support a non-positively curved metric, the geometrization conjecture states that such a 3-manifold contains a finite collection of pairwise disjoint, embedded 2-tori, and each component of the complement is either hyperbolic (supports a metric modeled on \mathbb{H}^3) or is non-positively curved Seifert fibered (supports a metric modeled on $\mathbb{H}^2 \times \mathbb{R}$). In the case where there are no hyperbolic components, the 3-manifold is an example of a *graph manifold*. The class of manifolds we consider are inspired by these notions.

DEFINITION 0.1. We will say that a compact smooth n -manifold M , $n \geq 3$, is a *graph manifold* provided that it can be constructed in the following way:

- (1) For every $i = 1, \dots, r$, take a complete finite-volume non-compact hyperbolic n_i -manifold N_i with toric cusps, where $3 \leq n_i \leq n$.
- (2) Denote by \bar{N}_i the manifold obtained by “truncating the cusps” of N_i , i.e. by removing from N_i a horospherical neighbourhood of each cusp.
- (3) Take the product $V_i = \bar{N}_i \times T^{n-n_i}$, where $T^k = (S^1)^k$ is the k -dimensional torus.
- (4) Fix a pairing of some boundary components of the V_i 's and glue the paired boundary components using *affine* diffeomorphisms of the boundary tori, so as to obtain a connected manifold of dimension n (see Section 2 for the precise definition of affine gluing in this context).

Observe that ∂M is either empty or consists of tori. The submanifolds V_1, \dots, V_r will be called the *pieces* of M . The manifold \bar{N}_i is the *base* of V_i , while every subset of the form $\{*\} \times T^{n-n_i} \subseteq V_i$ is a *fiber* of V_i . The boundary tori which are identified together will be called the *internal walls* of M (so any two distinct pieces in M will be separated by a collection of walls), while the components of ∂M will be called the *boundary walls* of M .

Informally, our manifolds can be decomposed into pieces, each of which supports a finite-volume product metric locally modeled on some $\mathbb{H}^k \times \mathbb{R}^{n-k}$ ($k \geq 3$). We

point out that our class of manifolds has some differences with the corresponding classical 3-dimensional graph manifolds. For instance:

- we do not allow pieces to be products of hyperbolic surfaces with codimension two tori (though see Chapter 13),
- we do not allow general finite volume quotients of $\mathbb{H}^k \times \mathbb{R}^{n-k}$, but
- we allow purely hyperbolic pieces in our decompositions (i.e. the case where $k = n$).

Our notion of generalized graph manifolds includes both the classical “double” of a finite volume hyperbolic manifold with toric cusps, as well as twisted doubles of such manifolds (see [ArFa]). We emphasize that, restricting down to 3-dimensions, our notion of graph manifold (and subsequent notion of extended graph manifold) **do not** coincide with the classical 3-dimensional graph manifolds.

Now our graph manifolds are “built up”, in a relatively simple manner, from non-positively curved manifolds. If we know some property holds for non-positively curved manifolds, and hence for all the pieces in our decomposition, we could expect it to hold for the graph manifold. This monograph pursues this general philosophy, with a view towards establishing analogues of various rigidity theorems for the class of graph manifolds. Chapter 1 starts out with a review of some basic notions: quasi-isometries, quasi-actions, and the Milnor-Sv̄arc Lemma.

In Chapter 2, we introduce our graph manifolds, and establish some basic general results. In Section 2.5, we provide a first family of examples of graph manifolds which *cannot* support any locally CAT(0)-metric. More precisely, we construct examples of graph manifolds M^n ($n \geq 5$) where the fundamental group of the walls is *not* quasi-isometrically embedded in $\pi_1(M)$. Such an M cannot support a locally CAT(0)-metric, by the Flat Torus theorem.

In Chapter 3, we study the topology of our graph manifolds. Recall that the *Borel Conjecture* states that if M, M' are aspherical manifolds with isomorphic fundamental group, then they are in fact homeomorphic. If the manifold M is assumed to support a Riemannian metric of non-positive curvature and has dimension ≥ 5 , then the validity of the Borel Conjecture is a celebrated result of Farrell-Jones. Our next result establishes (Section 3.3):

THEOREM 0.2 (Topological rigidity). *Let M be a graph manifold (possibly with boundary), of dimension $n \geq 6$. Assume M' is an arbitrary manifold and $\rho : M' \rightarrow M$ is a homotopy equivalence which restricts to a homeomorphism $\rho|_{\partial M'} : \partial M' \rightarrow \partial M$ between the boundaries of the manifolds. Then ρ is homotopic, rel ∂ , to a homeomorphism $\bar{\rho} : M' \rightarrow M$.*

Along the way, we show that our graph manifolds are always aspherical (Section 3.1) and show vanishing of the lower algebraic K -groups of the fundamental groups (Section 3.2). We also point out that the Baum-Connes conjecture holds (Section 3.4) and mention some well-known consequences. It is worth noting that, by work of Ontaneda [On, Theorem 1], there are examples of doubles of finite volume hyperbolic manifolds which support exotic PL-structures. As such, in Theorem 0.2, there are examples where no PL-homeomorphism (and hence, no diffeomorphism) exists between M and M' .

From the generalized Seifert-Van Kampen theorem, the fundamental group Γ of one of our graph manifolds M^n can be expressed as the fundamental group of a

graph of groups, with vertex groups given by the fundamental groups of the pieces, and edge groups isomorphic to \mathbb{Z}^{n-1} . To further develop our analysis of graph manifolds, we would like to ensure that reasonable maps between graph manifolds have to (essentially) preserve the pieces. The following result, which is the main goal of Chapter 4, is crucial:

THEOREM 0.3 (Isomorphisms preserve pieces). *Let M_1, M_2 be a pair of graph manifolds and let $\Gamma_i = \pi_1(M_i)$ be their respective fundamental groups. Let $\Lambda_1 \leq \Gamma_1$ be a subgroup conjugate to the fundamental group of piece V in M_1 , and $\phi : \Gamma_1 \rightarrow \Gamma_2$ be an isomorphism. Then $\phi(\Lambda_1)$ is conjugate to the fundamental group $\Lambda_2 \leq \Gamma_2$ of a piece in M_2 .*

A fairly straightforward consequence of this result is a necessary condition for two graph manifolds to have isomorphic fundamental groups (see also Theorem 4.4):

COROLLARY 0.4. *Let M, M' be a pair of graph manifolds. If $\phi : \pi_1(M) \rightarrow \pi_1(M')$ is an isomorphism, then it induces a graph isomorphism between the associated graph of groups. Moreover, vertices identified via this graph isomorphism must have associated vertex groups which are isomorphic.*

This result will be somewhat refined in Section 8.5.

In Chapter 5, we return to studying the topology of graph manifolds. Building on Theorem 0.3, we prove the following:

THEOREM 0.5 (Smooth rigidity). *Let M, M' be arbitrary graph manifolds, and let $\varphi : \pi_1(M) \rightarrow \pi_1(M')$ be a group isomorphism. Then φ is induced by a diffeomorphism $\psi : M \rightarrow M'$.*

Ontaneda [On] had previously shown smooth rigidity within the class of doubles of finite volume hyperbolic manifolds. Next, for M a smooth manifold, we denote by $\text{MCG}(M)$ the *mapping class group* of M , i.e. the group of homotopy classes of diffeomorphisms of M into itself. Theorem 0.5 easily implies the following corollary (see Section 5.4):

COROLLARY 0.6. *Let M be a graph manifold. Then, the group $\text{MCG}(M)$ is isomorphic to the group $\text{Out}(\pi_1(M))$ of the outer automorphisms of $\pi_1(M)$.*

Using Corollary 0.6, it is easy to see that $\text{MCG}(M)$ is often infinite. For example, this is always the case when considering doubles or twisted doubles (in the sense of [ArFa]) of one-cusped hyperbolic manifolds with toric cusp (see Remarks 5.2 and 5.6).

In Chapter 6, we shift perspective from topology to algebra, and describe some group theoretic properties of fundamental groups of graph manifolds. In order to properly state our results, we need to define the notion of *transverse gluings*.

DEFINITION 0.7. Let V_1, V_2 be a pair of adjacent pieces of the graph manifold M , and let $T^+ \subseteq \partial V_1, T^- \subseteq \partial V_2$ be two boundary tori identified by the pairing defining M . If $T \subseteq M$ is the torus corresponding to T^+ and T^- , we say that V_1 and V_2 have *transverse fibers* along T provided that, under the gluing diffeomorphism $\psi : T^+ \rightarrow T^-$, the image of the fiber subgroup of $\pi_1(T^+)$ under ψ_* intersects the fiber subgroup of $\pi_1(T^-)$ only in $\{0\}$. This is equivalent to asking that the sum of the dimensions of the fibers of T^+ and T^- is strictly less than the dimension of M , and that the image of every fiber of T^+ under ψ is transverse to every fiber of T^- .

Now putting together Propositions 6.1, 6.3, 6.9 and Theorem 6.5, we obtain the following:

THEOREM 0.8. *Let M be an arbitrary graph manifold.*

- (1) *If an arbitrary subgroup $H < \pi_1(M)$ has Kazhdan's property (T), then H is the trivial subgroup.*
- (2) *(Tits Alternative): If $H < \pi_1(M)$ is a finitely generated subgroup, then either H is solvable, or H contains a non-abelian free group.*
- (3) *Suppose that $\partial M = \emptyset$, and that M contains a pair of adjacent pieces with transverse fibers. Then $\pi_1(M)$ is co-Hopfian.*
- (4) *$\pi_1(M)$ has uniformly exponential growth.*

Starting from Chapter 7, we shift our focus to coarse geometric properties of our graph manifolds. As we mentioned earlier, there exist examples of graph manifolds M^n with the property that certain walls $T \subset M$ have fundamental groups $\pi_1(T) \cong \mathbb{Z}^{n-1} \hookrightarrow \pi_1(M^n)$ which are *not* quasi-isometrically embedded. As one might expect, the presence of such walls causes serious difficulties when trying to study the coarse geometry of M . In order to proceed further, we restrict the class of graph manifolds under consideration.

DEFINITION 0.9. A graph manifold is *irreducible* if every pair of adjacent pieces has transverse fibers along every common torus.

In the case of 1-dimensional fibers, a graph manifold is irreducible if and only if the S^1 -bundle structure on each piece cannot be extended to the union of adjacent pieces. Therefore, in the context of higher-dimensional graph manifolds, the irreducible ones probably provide the closest analogue of the classical 3-dimensional graph manifolds.

The key property of irreducible graph manifolds is that they have much better coarse geometric properties, as evidenced by our:

THEOREM 0.10. *Let M be an irreducible graph manifold. Then the fundamental group of every fiber, wall, and piece, is quasi-isometrically embedded in $\pi_1(M)$.*

The proof of this result occupies the bulk of Chapter 7 (see in particular Theorem 7.11 and Corollary 7.13).

In Chapter 8, we establish some group theoretic properties of the fundamental groups of irreducible graph manifolds. A very useful feature of irreducible graph manifolds is that the action of the fundamental group on the associated Bass-Serre tree is *acylindrical* (see Section 8.1). Putting together Corollary 8.3, Proposition 8.4 and Corollary 8.7, we establish the following:

THEOREM 0.11. *Let M be an irreducible graph manifold. Then*

- (1) *if M does not consist of a single piece with non-trivial fiber and no internal walls, then the group $\pi_1(M)$ is C^* -simple,*
- (2) *the group $\pi_1(M)$ has solvable word problem, and*
- (3) *if M has at least two internal walls, then $\pi_1(M)$ is SQ-universal.*

It is maybe worth mentioning that, in proving points (1) and (2) of the above Theorem, we establish some more general results about the C^* -simplicity and the solvability of the word problem for fundamental groups of graphs of groups (see Propositions 8.2 and 8.6).

Now recall that, by Corollary 0.4, to have any chance of having isomorphic fundamental groups, two graph manifolds would have to be built up using the exact same pieces, and the gluings would have to identify the same collection of boundary tori together. As such, the only possible variation lies in the choice of gluing maps used to identify the boundary tori together. In Section 8.5, we show how, in some cases, fixing the collection of pieces, we can still produce infinitely many non-isomorphic fundamental groups simply by varying the gluings between the common tori. The construction is flexible enough that we can even ensure that all the resulting graph manifolds are irreducible.

In Chapter 9, we start analyzing quasi-isometries between fundamental groups of irreducible graph manifolds. By studying the asymptotic cone of the universal cover of M , we are able to show:

THEOREM 0.12 (QI's preserve pieces of irreducible graph manifolds). *Let M_1, M_2 be a pair of irreducible graph manifolds, and $\Gamma_i = \pi_1(M_i)$ their respective fundamental groups. Let $\Lambda_1 \leq \Gamma_1$ be a subgroup conjugate to the fundamental group of a piece in M_1 , and $\varphi : \Gamma_1 \rightarrow \Gamma_2$ be a quasi-isometry. Then, the set $\varphi(\Lambda_1)$ is within finite Hausdorff distance from a conjugate of $\Lambda_2 \leq \Gamma_2$, where Λ_2 is the fundamental group of a piece in M_2 .*

Since pieces are essentially mapped to pieces under quasi-isometries, our next goal is to understand the behavior of groups quasi-isometric to the fundamental group of a piece. This is the subject of Chapter 10, where we establish:

THEOREM 0.13 (QI-rigidity of pieces). *Let N be a complete finite-volume hyperbolic m -manifold, $m \geq 3$, and let Γ be a finitely generated group quasi-isometric to $\pi_1(N) \times \mathbb{Z}^d$, $d \geq 0$. Then there exists a finite-index subgroup Γ' of Γ , a finite-sheeted covering N' of N , a group Δ and a finite group F such that the following short exact sequences hold:*

$$\begin{aligned} 1 &\longrightarrow \mathbb{Z}^d \xrightarrow{j} \Gamma' \longrightarrow \Delta \longrightarrow 1, \\ 1 &\longrightarrow F \longrightarrow \Delta \longrightarrow \pi_1(N') \longrightarrow 1. \end{aligned}$$

Moreover, $j(\mathbb{Z}^d)$ is contained in the center of Γ' . In other words, Γ' is a central extension by \mathbb{Z}^d of a finite extension of $\pi_1(N')$.

Note that the analogous result in the setting where N is compact has been established by Kleiner and Leeb [**KiLe**]. A consequence of this result is that we can determine when two pieces have quasi-isometric fundamental group: their fibers must be of the same dimension, while their bases must be commensurable.

In Chapter 11, we study groups quasi-isometric to an irreducible graph manifold, and show that they must exhibit a graph of groups structure which closely resembles that of a graph manifold (compare with [**MSW1**], [**MSW2**], [**Pa**]).

THEOREM 0.14. *Let M be an irreducible graph n -manifold obtained by gluing the pieces $V_i = \bar{N}_i \times T^{d_i}$, $i = 1, \dots, k$. Let Γ be a group quasi-isometric to $\pi_1(M)$. Then either Γ itself or a subgroup of Γ of index two is isomorphic to the fundamental group of a graph of groups satisfying the following conditions:*

- every edge group contains \mathbb{Z}^{n-1} as a subgroup of finite index;

- for every vertex group Γ_v there exist $i \in \{1, \dots, k\}$, a finite-sheeted covering N' of N_i and a finite-index subgroup Γ'_v of Γ_v that fits into the exact sequences

$$1 \longrightarrow \mathbb{Z}^{d_i} \xrightarrow{j} \Gamma'_v \longrightarrow \Delta \longrightarrow 1,$$

$$1 \longrightarrow F \longrightarrow \Delta \longrightarrow \pi_1(N') \longrightarrow 1,$$

where F is a finite group, and $j(\mathbb{Z}^{d_i})$ is contained in the center of Γ'_v .

As we mentioned at the beginning of this introduction, many of our rigidity results are inspired by corresponding results in the theory of non-positively curved spaces. In order to emphasize that our class of manifolds is genuinely different, we show in Chapter 12:

THEOREM 0.15. *In each dimension $n \geq 4$, there are infinitely many examples of n -dimensional irreducible graph manifolds which do not support any locally $CAT(0)$ -metric.*

In Chapter 13, we introduce the notion of an *extended graph manifold*. These are similar to our generalized graph manifolds, except that we now allow pieces of the form $\Sigma \times T^{n-2}$, where Σ is a finite volume hyperbolic surface. We also require the technical condition that in adjacent surface pieces, fibers are never identified together. In Theorem 13.3, we provide a characterization of the wall subgroups in extended graph n -manifolds. Using this, we can establish analogues of our main theorems for the class of extended graph n -manifolds: smooth rigidity (see Theorem 13.5) and the structure theorem for groups quasi-isometric to the fundamental group of such manifolds (see Theorem 13.8). Along the way, we consider groups quasi-isometric to the direct product $F_k \times \mathbb{Z}^d$ of a free group F_k with a free abelian group \mathbb{Z}^d . We show that any such group must contain a finite index subgroup which is isomorphic to some $F_{k'} \times \mathbb{Z}^d$.

Finally, in Chapter 14, we provide some concluding remarks, and propose various open problems suggested by our work.

Acknowledgments

The authors have benefitted from helpful conversations with a number of colleagues. We would particularly like to thank Cornelia Druțu, Tom Farrell, Max Forester, Mohamad Hindawi, Ian Leary, Peter Linnell, Ben McReynolds, Panos Papasoglu, Juan Souto, and Ralf Spatzier.

Lafont was partially supported by the NSF, under grant DMS-0906483, and by an Alfred P. Sloan research fellowship.

Part 1

Graph manifolds: topological and algebraic properties

CHAPTER 1

Quasi-isometries and quasi-actions

In this chapter we fix some notations we will extensively use in the rest of this monograph. We also list some well-known results about quasi-isometries and quasi-actions, providing a proof for the strengthened version of Milnor-Svarc's Lemma described in Lemma 1.4. Such a result is probably well-known to experts, but we did not find an appropriate reference for it in the literature.

Let $(X, d), (Y, d')$ be metric spaces and $k \geq 1, c \geq 0$ be real numbers. A (not necessarily continuous) map $f: X \rightarrow Y$ is a (k, c) -quasi-isometric embedding if for every $p, q \in X$ the following inequalities hold:

$$\frac{d(p, q)}{k} - c \leq d'(f(p), f(q)) \leq k \cdot d(p, q) + c.$$

Moreover, a (k, c) -quasi-isometric embedding f is a (k, c) -quasi-isometry if there exists a (k, c) -quasi-isometric embedding $g: Y \rightarrow X$ such that $d'(f(g(y)), y) \leq c, d(g(f(x)), x) \leq c$ for every $x \in X, y \in Y$. Such a map g is called a *quasi-inverse* of f . It is easily seen that a (k, c) -quasi-isometric embedding $f: X \rightarrow Y$ is a (k', c') -quasi-isometry for some $k' \geq 1, c' \geq 0$ if and only if its image is r -dense for some $r \geq 0$, *i.e.* if every point in Y is at distance at most r from some point in $f(X)$ (and in this case k', c' only depend on k, c, r).

1.1. The quasi-isometry type of a group

If Γ is a group endowed with a finite system of generators S such that $S = S^{-1}$, the *Cayley graph* $C_S(\Gamma)$ of Γ is the geodesic graph defined as follows: $C_S(\Gamma)$ has Γ as set of vertices, two vertices $g, g' \in C_S(\Gamma)$ are joined by an edge if and only if $g^{-1}g'$ lies in S , and every edge has unitary length. It is very easy to show that different finite sets of generators for the same group define quasi-isometric Cayley graphs, so every finitely generated group is endowed with a metric which is well-defined up to quasi-isometry.

REMARK 1.1. Suppose $i: \Gamma_1 \rightarrow \Gamma_2, j: \Gamma_2 \rightarrow \Gamma_3$ are injective group homomorphisms between finitely generated groups, and let S_i be a finite system of generators for $\Gamma_i, i = 1, 2, 3$. We may enlarge S_2 and S_3 in such a way that $i(S_1) \subseteq S_2, j(S_2) \subseteq S_3$. Under this assumption, both i and j are 1-Lipschitz embeddings with respect to the word metrics defined via the S_i 's. Using this fact, it is not hard to show that the composition $j \circ i$ is a quasi-isometric embedding if and only if both i and j are quasi-isometric embeddings.

1.2. The Milnor-Svarc Lemma

The following fundamental result shows how the quasi-isometry type of a group is related to the quasi-isometry type of a metric space on which the group acts geometrically. A geodesic metric space X is *proper* if every closed ball in X is

compact. An isometric action $\Gamma \times X \rightarrow X$ of a group Γ on a metric space X is *proper* if for every compact subset $K \subseteq X$ the set $\{g \in \Gamma \mid g \cdot K \cap K \neq \emptyset\}$ is finite, and *cocompact* if X/Γ is compact.

THEOREM 1.2 (Milnor-Svarc Lemma). *Suppose Γ acts by isometries, properly and cocompactly on a proper geodesic space X . Then Γ is finitely generated and quasi-isometric to X , a quasi-isometry being given by the map*

$$\psi: \Gamma \rightarrow X, \quad \psi(\gamma) = \gamma(x_0),$$

where $x_0 \in X$ is any basepoint.

As a corollary, if M is a compact Riemannian manifold with Riemannian universal covering \widetilde{M} , then the fundamental group of M is quasi-isometric to \widetilde{M} . A proof of this result can be found in [BrHa, Chapter I.8.19], and we will prove a slightly more general version of the Lemma in the next section.

1.3. From quasi-isometries to quasi-actions

Suppose (X, d) is a geodesic metric space, let $\text{QI}(X)$ be the set of quasi-isometries of X into itself, and let Γ be a group. For $k \geq 1$, a *k-quasi-action* of Γ on X is a map $h: \Gamma \rightarrow \text{QI}(X)$ such that the following conditions hold:

- (1) $h(\gamma)$ is a (k, k) -quasi-isometry with k -dense image for every $\gamma \in \Gamma$;
- (2) $d(h(1)(x), x) \leq k$ for every $x \in X$;
- (3) the composition $h(\gamma_1) \circ h(\gamma_2)$ is at distance bounded by k from the quasi-isometry $h(\gamma_1\gamma_2)$, i.e.

$$d(h(\gamma_1\gamma_2)(x), h(\gamma_1)(h(\gamma_2)(x))) \leq k \quad \text{for every } x \in X, \gamma_1, \gamma_2 \in \Gamma.$$

A k -quasi-action h as above is *k'-cobounded* if every orbit of Γ in X is k' -dense. A (cobounded) quasi-action is a map which is a (k' -cobounded) k -quasi-action for some $k, k' \geq 1$. Throughout the whole paper, by an abuse of notation, when h is a quasi-action as above we do not distinguish between γ and $h(\gamma)$.

REMARK 1.3. If h is a k -quasi-action as above, then for every $\gamma \in \Gamma, x_0, x_1, p \in X$ we have

$$d(\gamma(x_1), p) \leq d(\gamma(x_1), \gamma(x_0)) + d(\gamma(x_0), p) \leq kd(x_0, x_1) + k + d(\gamma(x_0), p).$$

Using this inequality, it is not difficult to show that if there exists a k' -dense orbit of Γ in X , then h is k'' -cobounded for some k'' (possibly larger than k').

Suppose M is a geodesic metric space with metric universal covering \widetilde{M} , let Γ be a finitely generated group and suppose we are given a quasi-isometry $\tilde{\varphi}: \Gamma \rightarrow \pi_1(M)$. We now briefly recall the well-known fact that $\tilde{\varphi}$ naturally induces a cobounded quasi-action of Γ on \widetilde{M} .

Let $\varphi: \Gamma \rightarrow \widetilde{M}$ be a fixed quasi-isometry provided by Milnor-Svarc's Lemma, and let $\psi: \widetilde{M} \rightarrow \Gamma$ be a quasi-inverse of φ . For each $\gamma \in \Gamma$ we define a map $h(\gamma): \widetilde{M} \rightarrow \widetilde{M}$ by setting

$$h(\gamma)(x) = \varphi(\gamma \cdot \psi(x)) \quad \text{for every } x \in \widetilde{M}.$$

Since $h(1) = \varphi \circ \psi$, the map $h(1)$ is at finite distance from the identity of \widetilde{M} . The left multiplication by a fixed element of Γ defines an isometry of any Cayley graph of Γ , so each $h(\gamma)$ is the composition of three quasi-isometries with fixed constants. In particular, it is a quasi-isometry and its quasi-isometry constants can be bounded by

a universal constant which only depends on φ and ψ , and is therefore independent of γ . As such, we have that for every $\gamma \in \Gamma$ the map $h(\gamma)$ is a (k, k) -quasi-isometry with k -dense image, where k is some fixed uniform constant. Moreover, it is easily seen that for each γ_1, γ_2 , $h(\gamma_1\gamma_2)$ is at a finite distance (bounded independently of γ_1, γ_2) from $h(\gamma_1) \circ h(\gamma_2)$, that is, h defines a quasi-action. Since every Γ -orbit in Γ is 1-dense, the quasi-action h is clearly cobounded.

In Chapters 10 and 11 we need the following strengthened version of Milnor-Svarc's Lemma.

LEMMA 1.4. *Let X be a geodesic space with basepoint x_0 , and let Γ be a group. Let $h: \Gamma \rightarrow \text{QI}(X)$ be a cobounded quasi-action of Γ on X , and suppose that for each $r > 0$, the set $\{\gamma \in \Gamma \mid \gamma(B(x_0, r)) \cap B(x_0, r) \neq \emptyset\}$ is finite. Then Γ is finitely generated and the map $\varphi: \Gamma \rightarrow X$ defined by $\varphi(\gamma) = \gamma(x_0)$ is a quasi-isometry.*

PROOF. The usual proof of Milnor-Svarc's Lemma works in this case too, up to minor changes. We will closely follow [BrHa, Chapter I.8.19]. Suppose that h is a k -cobounded k -quasi-action, and let us first prove that the finite set

$$\mathcal{A} = \{\gamma \in \Gamma \mid \gamma(B(x_0, 2k^2 + 5k)) \cap B(x_0, 2k^2 + 5k) \neq \emptyset\}$$

generates Γ . Fix $\gamma \in \Gamma$ and consider a geodesic $\alpha: [0, 1] \rightarrow X$ joining x_0 with $\gamma(x_0)$. If $n \in \mathbb{N}$ is such that $d(x_0, \gamma(x_0)) \leq n \leq d(x_0, \gamma(x_0)) + 1$, we can choose $0 = t_0 < \dots < t_n = 1$ in such a way $d(\alpha(t_i), \alpha(t_{i+1})) \leq 1$ for each i . For each t_i pick γ_i so that $d(\alpha(t_i), \gamma_i(x_0)) \leq k$, with $\gamma_0 = 1$ and $\gamma_n = \gamma$, and observe that $d(\gamma_i(x_0), \gamma_{i+1}(x_0)) \leq 2k + 1$ for $i = 0, \dots, n-1$. Since

$$\begin{aligned} d(x_0, (\gamma_i^{-1}\gamma_{i+1})(x_0)) &\leq d(\gamma_i^{-1}(\gamma_i(x_0)), \gamma_i^{-1}(\gamma_{i+1}(x_0))) + 3k \\ &\leq kd(\gamma_i(x_0), \gamma_{i+1}(x_0)) + 4k \\ &\leq k(2k + 1) + 4k \end{aligned}$$

we see that $\gamma_i^{-1}\gamma_{i+1} \in \mathcal{A}$. This tells us that

$$\gamma = \gamma_0(\gamma_0^{-1}\gamma_1) \dots (\gamma_{n-1}^{-1}\gamma_n)$$

is a product of at most $d(x_0, \gamma(x_0)) + 1$ elements of \mathcal{A} . But γ was chosen arbitrarily, so \mathcal{A} is indeed a generating set for Γ .

Moreover, if $d_{\mathcal{A}}$ is the word metric with respect to \mathcal{A} , we have $d_{\mathcal{A}}(1, \gamma) \leq d(x_0, \gamma(x_0)) + 1$, and for every $\gamma, \gamma' \in \Gamma$ we have

$$\begin{aligned} d_{\mathcal{A}}(\gamma, \gamma') &= d_{\mathcal{A}}(1, \gamma^{-1}\gamma') \leq d(x_0, (\gamma^{-1}\gamma')(x_0)) + 1 \\ &\leq d(\gamma^{-1}(\gamma(x_0)), \gamma^{-1}(\gamma'(x_0))) + 3k + 1 \\ &\leq kd(\gamma(x_0), \gamma'(x_0)) + 4k + 1 \end{aligned}$$

which is one of the two inequalities needed to prove that φ is a quasi-isometric embedding. For the reverse inequality, we first establish a useful inequality. For an arbitrary pair of elements γ_1, γ_2 in Γ , we have the estimate:

$$\begin{aligned} d(\gamma_1(x_0), \gamma_2(x_0)) &= d(\gamma_1(x_0), (\gamma_1\gamma_1^{-1})(\gamma_2(x_0))) + k \\ &\leq d(\gamma_1(x_0), \gamma_1(\gamma_1^{-1}(\gamma_2(x_0)))) + 2k \\ &\leq kd(x_0, \gamma_1^{-1}(\gamma_2(x_0))) + 3k \\ &\leq kd(x_0, (\gamma_1^{-1}\gamma_2)(x_0)) + k^2 + 3k \end{aligned}$$

Choose μ so that $d(x_0, a(x_0)) \leq \mu$ for each $a \in \mathcal{A}$. Given any two elements $\gamma, \gamma' \in \Gamma$, let $n = d_{\mathcal{A}}(\gamma, \gamma')$ and write $\gamma^{-1}\gamma' = a_1 \dots a_n$, where $a_i \in \mathcal{A}$. Set $g_0 = 1$, $g_i = a_1 \dots a_i$, $i = 1, \dots, n$, so that $g_n = \gamma^{-1}\gamma'$. From the above inequality, we see that $d(g_i(x_0), g_{i+1}(x_0)) \leq k\mu + k^2 + 3k$ for every $i = 0, \dots, n-1$. Combining this estimate with the above inequality, we finally obtain

$$\begin{aligned} d(\gamma(x_0), \gamma'(x_0)) &\leq kd(x_0, g_n(x_0)) + k^2 + 3k \\ &\leq k\left(\sum_{i=1}^n d(g_{i-1}(x_0), g_i(x_0))\right) + k^2 + 3k \\ &\leq k(k\mu + k^2 + 3k)d_{\mathcal{A}}(\gamma, \gamma') + k^2 + 3k. \end{aligned}$$

We have thus proved that φ is a quasi-isometric embedding, and the fact that h is cobounded now implies that it is in fact a quasi-isometry. \square

Generalized graph manifolds

Let us introduce the precise definition of high dimensional graph manifold. Fix $n \geq 3$, $k \in \mathbb{N}$ and $n_i \in \mathbb{N}$ with $3 \leq n_i \leq n$, and for every $i = 1, \dots, k$ let N_i be a complete finite-volume non-compact hyperbolic n_i -manifold with toric cusps. It is well-known that each cusp of N_i supports a canonical smooth foliation by closed tori, which defines in turn a diffeomorphism between the cusp and $T^{n_i-1} \times [0, \infty)$, where $T^{n_i-1} = \mathbb{R}^{n_i-1}/\mathbb{Z}^{n_i-1}$ is the standard torus. Moreover, the restriction of the hyperbolic metric to each leaf of the foliation induces a flat metric on each torus, and there is a canonical affine diffeomorphism between any such two leaves.

We now “truncate” the cusps of N_i by setting $\overline{N}_i = N_i \setminus \cup_{j=1}^{a_i} T_j^{n_i-1} \times (4, \infty)$, where $T_j^{n_i-1} \times [0, \infty)$, $j = 1, \dots, a_i$ are the cusps of V_i . If $V_i = \overline{N}_i \times T^{n-n_i}$, then V_i is a well-defined smooth manifold with boundary, and as mentioned above the boundary of V_i is endowed with a well-defined affine structure. Moreover, the boundary of V_i admits a collar which is canonically foliated by affine tori.

Let now \mathcal{B} be a subset of the set of boundary components of the V_i 's, and suppose that a pairing of the boundary components in \mathcal{B} is fixed. We can construct a smooth manifold M by gluing the V_i 's along affine diffeomorphisms between the paired tori in \mathcal{B} : the smooth manifold M obtained in this way is what we call a *graph n -manifold*. The manifolds V_1, \dots, V_k (which will be often considered as subsets of M itself) are called the *pieces* of M . For every i , we say that N_i (or \overline{N}_i) is the *base* of V_i , while if $p \in \overline{N}_i$, then the set $\{p\} \times T^{n-n_i} \subseteq V_i$ is a *fiber* of V_i . Abusing terminology, we will sometimes also refer to T^{n-n_i} as the fiber of V_i .

Observe that M is closed (*i.e.* $\partial M = \emptyset$) if and only if \mathcal{B} coincides with the whole set of boundary components of the V_i 's.

REMARK 2.1. The product of an affine torus with a truncated hyperbolic manifold with toric cusps provides the simplest example of graph manifold with non-empty boundary. The quasi-isometry type of the fundamental group of such a manifold will be studied in detail in Chapter 10.

REMARK 2.2. The simplest examples of closed graph manifolds are obtained just by gluing (truncated) hyperbolic manifolds with toric cusps along affine diffeomorphisms of their horospherical boundaries. Therefore, it makes sense to compare our rigidity results with the analogous results described in [On] (for doubles of cusped hyperbolic manifolds), in [ArFa] (for twisted doubles of cusped hyperbolic manifolds), and in [Ng] (for manifolds obtained by gluing locally symmetric negatively curved manifolds with deleted cusps).

REMARK 2.3. Suppose \overline{N} and \overline{N}' are obtained as above by deleting from a hyperbolic N horospherical cusp neighbourhoods of possibly different “heights”. Then, there exists a diffeomorphism between \overline{N} and \overline{N}' which is coherent with the

identification of $\partial\bar{N}$ and $\partial\bar{N}'$ induced by the canonical foliations of the cusps of N . In particular, the diffeomorphism type of a graph manifold M does not depend on the choice of the height of the cusps removed from the hyperbolic factors of the pieces into which M decomposes.

REMARK 2.4. It is proved in [HsWa] that, if $n \geq 5$, then any diffeomorphism between affine n -dimensional tori is C^0 -isotopic to an affine diffeomorphism. As a consequence, for $n \geq 6$, if we allow also non-affine gluings, then we do not obtain new homeomorphism types of graph manifolds. On the other hand, as showed in [ArFa], requiring the gluings to be affine is necessary for getting smooth rigidity results as in our Theorem 0.5 (i.e. non-affine gluings can give rise to new diffeomorphism types of manifolds).

2.1. Putting a metric on a high dimensional graph manifold

By construction, each hypersurface in M corresponding to a boundary torus of some V_i is either a boundary component of M , or admits a canonical smooth bicollar in M diffeomorphic to $T^{n-1} \times [-3, 3]$, which is obtained by gluing, according to the pairing of the boundary components in \mathcal{B} , some subsets of the form $\partial V_i \times [1, 4]$, where ∂V_i is canonically identified with $\partial V_i \times \{4\}$.

In what follows, we will say that a point $p \in T^{n-1} \times \{-3\}$ is *tied* to $q \in T^{n-1} \times \{3\}$ if $p = (x, -3)$, $q = (x, 3)$ for some $x \in T^{n-1}$, i.e. if p, q have the same ‘toric’ component in the product space $T^{n-1} \times [-3, 3] \subseteq M$.

The following lemma shows how one can put on M a Riemannian metric which somewhat extends the product metrics defined on the V_i 's.

LEMMA 2.5. *Consider $A_1 = T^k \times [-3, 0]$ and $A_2 = T^k \times [0, 3]$, each equipped with a Riemannian metric g_i , and let $B_1 = T^k \times [-3, -2]$, $B_2 = T^k \times [2, 3]$. Then there exists a Riemannian metric on $A = T^k \times [-3, 3]$ such that $g|_{B_i} = g_i|_{B_i}$, $i = 1, 2$.*

PROOF. Let $\rho : [-3, 3] \rightarrow [-3, 3]$ be an odd C^∞ function such that:

- (1) $\rho|_{[2, 3]} = id$,
- (2) $\rho([1, 2]) = [0, 2]$,
- (3) $\rho|_{[0, 1]} = 0$.

Also, let $\delta : [-1/2, 1/2] \rightarrow [0, 1]$ be an increasing C^∞ function which is constantly 0 (resp. 1) in a neighborhood of $-1/2$ (resp. $1/2$) and is strictly positive in $[0, 1/2]$. We can define g as follows:

$$g(p, x) = \begin{cases} g_1(p, \rho(x)) & \text{for } x \in [-3, -1/2] \\ \delta(-x)g_1(p, 0) + \delta(x)g_2(p, 0) & \text{for } x \in [-1/2, 1/2] \\ g_2(p, \rho(x)) & \text{for } x \in [1/2, 3] \end{cases}$$

for all $p \in T^k, x \in [-3, 3]$. □

DEFINITION 2.6. We recall that a metric space X is *geodesic* if for every $x, y \in X$ there exists a rectifiable curve $\gamma : [0, 1] \rightarrow X$ joining x to y whose length is equal to $d(x, y)$ (the constant speed parameterization of such a curve is called *geodesic*). Suppose S is a submanifold of the (possibly bounded) simply connected Riemannian manifold X , and let d be the Riemannian metric of X . We say that S is *totally geodesic* in (X, d) (in the metric sense) if for every $p, q \in S$ there exists a geodesic of X which joins p to q and whose support is contained in S . In this case, the path

metric associated to the restriction of d to S coincides with the restriction of d to S .

Suppose \overline{N} is obtained by deleting some horospherical cusp neighbourhoods from a complete finite-volume hyperbolic manifold N , and let $\widetilde{\overline{N}}$ be the universal covering of \overline{N} , endowed with the metric induced by the covering map. Then, it is well-known (see *e.g.* [BrHa, pgs. 362-366]) that every component of $\partial\widetilde{\overline{N}}$ is totally geodesic in $\widetilde{\overline{N}}$, even if its extrinsic curvature in $\widetilde{\overline{N}}$ does not vanish. Putting together this observation with Lemma 2.5 we get the following:

COROLLARY 2.7. *Suppose M is a graph manifold, and let $U \subseteq M$ be the union of the bicollars of the toric hypersurfaces of $M \setminus \partial M$ corresponding to the boundary components of the pieces of M . Then M admits a Riemannian metric g which extends the restriction to $M \setminus U$ of the product metrics originally defined on the pieces of M .*

2.2. $\pi_1(M)$ as the fundamental group of a graph of groups

The decomposition of a graph n -manifold M into pieces V_1, \dots, V_k induces on $\pi_1(M)$ the structure of the fundamental group of a graph of groups \mathcal{G}_M (see [Se] for the definition and some basic results on the fundamental group of a graph of groups). More precisely, let \mathcal{G}_M be the graph of groups that describes the decomposition of M into the V_i 's, in such a way that every vertex group is the fundamental group of the corresponding piece, every edge group is isomorphic to \mathbb{Z}^{n-1} , and the homomorphism of every edge group into the group of an adjacent vertex is induced by the inclusion of the corresponding boundary component of V_i into V_i . Then we have an isomorphism $\pi_1(M) \cong \pi_1(\mathcal{G}_M)$.

Recall that cusps of hyperbolic manifolds are π_1 -injective, so every boundary component of V_i is π_1 -injective in V_i . This implies that every piece (hence every boundary component of a piece) is π_1 -injective in M .

For later reference, we point out the following lemma, which can be easily deduced from [BePe, Lemma D.2.3]:

LEMMA 2.8. *Let N be a complete finite-volume hyperbolic n -manifold, $n \geq 3$.*

- (1) *Suppose that the cusps of N are toric, and that γ is a non-trivial element of $\pi_1(N)$. Then, the centralizer of γ in $\pi_1(N)$ is free abelian.*
- (2) *The center of $\pi_1(N)$ is trivial.*

The following remark is an immediate consequence of Lemma 2.8-(2).

REMARK 2.9. If N is a complete finite-volume hyperbolic n -manifold and d is a natural number, then the center of $\pi_1(N) \times \mathbb{Z}^d$ is given by $\{1\} \times \mathbb{Z}^d$. Therefore, if $V_i \cong \overline{N}_i \times T^d$ is a piece of M and $p_i: V_i \rightarrow N_i$ is the natural projection, then the center of $\pi_1(V_i)$ coincides with $\ker(p_i)_*$.

DEFINITION 2.10. Let V_i be a piece of M . Then the center of $\pi_1(V_i)$ is called the *fiber subgroup* of $\pi_1(V_i)$. If T is a component of ∂V_i , we call *fiber subgroup* of $\pi_1(T)$ the intersection of $\pi_1(T)$ with the fiber subgroup of $\pi_1(V_i)$.

2.3. The universal cover of M as a tree of spaces

In this subsection we begin our analysis of the metric structure of the universal covering \widetilde{M} of M . We will be mainly interested in the study of the quasi-isometric properties of \widetilde{M} .

DEFINITION 2.11. A *tree of spaces* (X, p, T) is a topological space X equipped with a map p on a (simplicial, but possibly not locally finite) tree T with the following property: for any edge e in T and t in the internal part e° of e , if $X_e = p^{-1}(t)$ then $p^{-1}(e^\circ)$ is homeomorphic to $X_e \times (0, 1)$.

DEFINITION 2.12. Suppose (X, p, T) is a tree of spaces where X is a Riemannian manifold. An *internal wall* of X is the closure of the preimage under p of the interior of an edge of T ; a *boundary wall* of X is simply a connected component of ∂X . If W is a (boundary or internal) wall of X , we will denote by d_W the path metric induced on W by the restriction to W of the Riemannian structure of X . A *chamber* $C \subseteq X$ is the preimage under p of a vertex of T ; we will denote by d_C the path metric induced on C by the restriction to C of the Riemannian structure of X . Two distinct chambers of X are *adjacent* if the corresponding vertices of T are joined by an edge, while a wall W is *adjacent* to the chamber C if $W \cap C \neq \emptyset$ (if W is internal, then W is adjacent to C if and only if the vertex corresponding to C is an endpoint of the edge corresponding to W , while if W is a boundary wall, then W is adjacent to C if and only if $W \subseteq C$).

Let us now come back to our graph n -manifold M . If $\dim N_i = n_i$, the universal covering of \overline{N}_i is isometric (as a Riemannian manifold) to the complement B_i in \mathbb{H}^{n_i} of an equivariant family of open disjoint horoballs. Following Schwartz, we say that B_i is a *neutered space*. In the rest of this monograph, we will extensively use several features of neutered spaces (see for example Proposition 7.4 or Section 9.3, where we will deduce asymptotic properties of such spaces from the well-know fact that they are *relatively hyperbolic* in the metric sense).

Since the fundamental group of each \overline{N}_i and each V_i injects in the fundamental group of $\pi_1(M)$, the universal coverings $\widetilde{V}_i = B_i \times \mathbb{R}^{n-n_i}$ embed into \widetilde{M} . Putting together this observation and Corollary 2.7 we get the following:

COROLLARY 2.13. M admits a Riemannian metric such that \widetilde{M} can be turned into a tree of spaces such that:

- (1) If C is a chamber of \widetilde{M} , then (C, d_C) is isometric (as a Riemannian manifold) to $B \times \mathbb{R}^k$, where B is a neutered space in \mathbb{H}^{n-k} .
- (2) If W is an internal wall of \widetilde{M} , then W is diffeomorphic to $\mathbb{R}^{n-1} \times [-1, 1]$.
- (3) If W is a boundary wall of \widetilde{M} , then W is isometric (as a Riemannian manifold) to \mathbb{R}^{n-1} .

We will call B the *base* of C , and $F = \mathbb{R}^k$ the *fiber* of C . If $\pi_B: C \rightarrow B$, $\pi_F: C \rightarrow \mathbb{R}^k$ are the natural projections, we will abuse the terminology, and also refer to a subset $F \subseteq C$ of the form $F = \pi_B^{-1}(x_0)$, where x_0 is a point in B , as a *fiber* of C . A fiber of \widetilde{M} is a fiber of some chamber of \widetilde{M} .

If $x, y \in C$, we denote by $d_B(x, y)$ the distance (with respect to the path metric of B) between $\pi_B(x)$ and $\pi_B(y)$, and by $d_F(x, y)$ the distance between $\pi_F(x)$ and $\pi_F(y)$ (so by construction $d_C^2 = d_B^2 + d_F^2$).

If (\widetilde{M}, p, T) is the tree of spaces described in Corollary 2.13, we refer to T as to the *Bass-Serre tree* of $\pi_1(M)$ (with respect to the isomorphism $\pi_1(M) \cong \pi_1(\mathcal{G}_M)$, or to the decomposition of M into the V_i 's). The action of $\pi_1(M)$ on \widetilde{M} induces an action of $\pi_1(M)$ on T . By the very definitions, (every conjugate of) the fundamental group of a piece (resp. of a paired boundary component of a piece) coincides with the stabilizer of a vertex (resp. of an edge) of T , and vice versa.

LEMMA 2.14. *Set $G = \pi_1(M)$. Let V_1, V_2 be pieces of M and T_i a component of ∂V_i , $i = 1, 2$. Let $G_i < \pi_1(M)$ (resp. $H_i < \pi_1(M)$) be (any conjugate of) the fundamental group of V_i (resp. of T_i). Then:*

- (1) *The normalizer of H_1 in G is equal to H_1 .*
- (2) *If H_2 is conjugate to H_1 in G , then $T_1 = T_2$.*
- (3) *The normalizer of G_1 in G is equal to G_1 .*
- (4) *If G_1 is conjugate to G_2 in G , then $V_1 = V_2$.*
- (5) *If $g \in G$ is such that $G_1 \cap gG_1g^{-1} \supseteq H_1$, then either $g \in G_1$ or V_1 is glued to itself along T_1 in M .*

PROOF. Let us consider the action of G on the Bass-Serre tree T corresponding to the decomposition of M into pieces.

(1): The subgroup H_1 coincides with the stabilizer of an edge $e_1 \subseteq T$. We first show that e_1 is the unique edge fixed by the whole H_1 . In fact, suppose that H_1 fixes the edge $e'_1 \neq e_1$ of T . Then H_1 acts trivially on the path joining e_1 to e'_1 . In particular, there exists an edge $e''_1 \subseteq T$ adjacent to e_1 which is fixed by H_1 . Let now $V = \overline{N} \times T^a$ be the piece of M corresponding to the vertex $v = e_1 \cap e''_1$, and let B be the neutered space covering \overline{N} . Up to replacing H_1 with a conjugate subgroup, we may suppose that $H_1 \subseteq \text{stab}(v) = \pi_1(V)$. Moreover, if an element $(g, w) \in \pi_1(V) = \pi_1(\overline{N}) \times \pi_1(T^a)$ fixes both e_1 and e''_1 , then the action of g on B has to leave invariant two horospherical components of ∂B . This implies that g has to fix the geodesic connecting such components, and since non-trivial covering automorphisms cannot have fixed points this gives in turn $g = 1$. We have thus proved that an element of G which stabilizes both e_1 and e''_1 has to belong to the fiber subgroup of the stabilizer of $v = e_1 \cap e''_1$. The fact that the rank of this fiber subgroup is strictly less than the rank of H_1 implies that H_1 cannot fix both e_1 and e''_1 , so it cannot fix e'_1 .

Suppose now that $g \in G$ is such that $gH_1g^{-1} = H_1$. Then H_1 is contained in the stabilizer of $g(e_1)$, so $g(e_1) = e_1$, and $g \in H_1$, as claimed in (1).

(2): As just shown in the proof of (1), the subgroups H_1, H_2 uniquely determine edges e_1, e_2 of T such that e_i is fixed by H_i , $i = 1, 2$. Suppose now that there exists $g \in G$ such that $gH_1g^{-1} = H_2$. Then H_2 fixes both e_2 and $g(e_1)$, so $e_2 = g(e_1)$. By the very construction of T , this implies that the covering automorphism $g: \widetilde{M} \rightarrow \widetilde{M}$ sends a wall covering T_1 onto a wall covering T_2 , so $T_1 = T_2$.

(3): Let v_1 be a vertex fixed by G_1 . As in the proof of (1), it is sufficient to show that v_1 is the only vertex fixed by G_1 . However, if G_1 fixes another vertex $v'_1 \neq v_1$, then it fixes an edge e exiting from v_1 . This implies that G_1 is contained in the stabilizer of an edge, which is clearly impossible since edge stabilizers are abelian.

(4): Let v_1, v_2 be the vertices of T fixed respectively by G_1, G_2 (v, v' are uniquely determined - see the proof of (3)), and suppose that there exists $g \in G$ such that $gG_1g^{-1} = G_2$. Then G_1 fixes both v_2 and $g(v_1)$, so $v_2 = g(v_1)$ by the

proof of (3). Therefore, the covering automorphism $g: \widetilde{M} \rightarrow \widetilde{M}$ sends a chamber covering V_1 onto a chamber covering V_2 , so $V_1 = V_2$.

(5): Let v_1, v'_1 be the vertices of T associated to G_1, gG_1g^{-1} respectively, and let e_1 be the edge fixed by H_1 , so that $v'_1 = g(v_1)$. If $v'_1 = v_1$, then g belongs to the stabilizer of v_1 , which is G_1 , and we are done. Otherwise, the assumption $G_1 \cap gG_1g^{-1} \supseteq H_1$ implies that every element of H_1 fixes v_1, v'_1 and e . Therefore, H_1 fixes the path in T joining v_1 with v'_1 . As showed in the proof of (1), however, H_1 fixes exactly one edge of T , and this implies that e_1 joins v_1 to $v'_1 = g(v_1)$, whence the conclusion. \square

2.4. Basic metric properties of \widetilde{M}

In this subsection we collect several metric properties of \widetilde{M} that we will extensively use in the following chapters in order to study the quasi-isometry type of the fundamental group of a graph manifold.

Recall from Corollary 2.13 that, if C is a chamber of \widetilde{M} , then (C, d_C) is isometric to the product of a neutered space with a Euclidean space. An elementary application of Milnor-Svarc Lemma (see Theorem 1.2) implies the following:

LEMMA 2.15. *If W is a wall of \widetilde{M} , then (W, d_W) is quasi-isometric to \mathbb{R}^{n-1} .*

Also recall that d denotes the distance associated to the Riemannian structure of \widetilde{M} . For every $r \geq 0$ and $X \subseteq \widetilde{M}$, we denote by $N_r(X) \subseteq \widetilde{M}$ the r -neighbourhood of X in \widetilde{M} , with respect to the metric d .

LEMMA 2.16. *If C is a chamber of \widetilde{M} , then there exists a function $g: \mathbb{R}^+ \rightarrow \mathbb{R}^+$ such that $g(t)$ tends to $+\infty$ as t tends to $+\infty$ and $d(x, y) \geq g(d_C(x, y))$ for each $x, y \in C$.*

PROOF. By quasi-homogeneity of C it is enough to prove the statement for a fixed x . Let us observe that d and d_C induce the same topology on C . Take any sequence $\{y_i\}$ of points such that $d_C(x, y_i)$ tends to $+\infty$. Since \widetilde{M} is proper, if the $d(x, y_i)$'s are bounded, then up to passing to a subsequence we can suppose $\lim_{i \rightarrow \infty} y_i = y$ for some $y \in \widetilde{M}$. But C is closed in \widetilde{M} , so we have $y \in C$. It is easily seen that this contradicts $d_C(x, y_i) \rightarrow +\infty$. \square

LEMMA 2.17. *Let W_1, W_2 be walls of \widetilde{M} , and suppose that there exists $r \in \mathbb{R}^+$ such that $W_1 \subseteq N_r(W_2)$. Then $W_1 = W_2$. In particular, distinct walls of \widetilde{M} lie at infinite Hausdorff distance from each other.*

PROOF. Considering the realization of \widetilde{M} as a tree of spaces, one can easily reduce to the case that W_1 and W_2 are adjacent to the same chamber C . By Lemma 2.16, up to increasing r we may assume that W_1 is contained in the r -neighbourhood of W_2 with respect to the path distance d_C of C .

Let $C = B \times \mathbb{R}^k$ be the decomposition of C into the product of a neutered space and a Euclidean space. Then, W_1 and W_2 project onto two horospheres O_1, O_2 of $B \subseteq \mathbb{H}^{n-k}$, and O_1 is contained in the r -neighbourhood of O_2 with respect to the distance d_B . Now, the distance d_B is bounded below by the restriction of the hyperbolic distance $d_{\mathbb{H}}$ of \mathbb{H}^{n-k} , so O_1 is contained in the r -neighbourhood of O_2 with respect to the distance $d_{\mathbb{H}}$. This forces $O_1 = O_2$, whence $W_1 = W_2$. \square

COROLLARY 2.18. *Let W (resp. C_1, C_2) be a wall (resp. two chambers) of \widetilde{M} . Then:*

- (1) *if $W \subseteq N_r(C_1)$ for some $r \geq 0$, then W is adjacent to C_1 ;*
- (2) *if $C_1 \subseteq N_r(C_2)$ for some $r \geq 0$, then $C_1 = C_2$; in particular, the Hausdorff distance between distinct chambers of \widetilde{M} is infinite.*

PROOF. (1) By considering the realization of \widetilde{M} as a tree of spaces, it is immediate to realize that W is contained in the r -neighbourhood of a wall adjacent to C_1 , so W is adjacent to C_1 by Lemma 2.17.

(2) Suppose W, W' are distinct walls both adjacent to C_1 . Then, by point (1) they are adjacent also to C_2 , and this forces $C_1 = C_2$. \square

In order to study the quasi-isometry type of \widetilde{M} , it would be very useful to know that the inclusions of walls and chambers are quasi-isometric embeddings. However, this is not true in general, as it is shown in the proof of Proposition 2.19 below, where we exploit this fact for constructing graph manifolds which do not support any CAT(0) metric.

In Chapter 7 we will define the class of *irreducible* graph manifolds, and we will prove that walls and chambers are quasi-isometrically embedded in the universal covering of an irreducible graph manifold.

2.5. Examples not supporting any locally CAT(0) metric

In this section we construct graph manifolds which do not support any locally CAT(0) metric. The construction described here is easy, and it is based on a straightforward application of the Flat Torus Theorem (see *e.g.* [BrHa, Chapter II.7]). As mentioned in the Introduction, however, there are reasons for being interested in *irreducible* graph manifolds (see Chapter 7). It turns out that providing examples of irreducible graph manifolds which do not support any locally CAT(0) metric is much harder. We will discuss this issue in detail in Chapter 12.

PROPOSITION 2.19. *Let $n \geq 3$, and take a hyperbolic n -manifold N with at least two cusps. We suppose as usual that every cusp of N is toric. For $i = 1, 2$, let $N_i = N$ and $V_i = \overline{N}_i \times T^2$. Then, we can glue the pieces V_1 and V_2 in such a way that the resulting graph manifold M does not support any CAT(0) metric.*

PROOF. Let A, A' be two distinct boundary tori of \overline{N} , and let $A_i \times T^2, A'_i \times T^2$ be the corresponding boundary tori of V_i . We now glue V_1 to V_2 as follows: $A_1 \times T^2$ is glued to $A_2 \times T^2$ with the identity, where A_1, A_2 are identified with A ; $A'_1 \times T^2$ is glued to $A'_2 \times T^2$ by an affine map φ such that $\varphi_*: \pi_1(A'_1 \times T^2) \rightarrow \pi_1(A'_2 \times T^2)$ is given by $\varphi_*(\bar{a}, c, d) = (\bar{a}, c, c + d)$, where $\bar{a} \in \mathbb{Z}^{n-1}$ and we are identifying A'_i with A' , and $\pi_1(A'_i \times T^2) = \pi_1(A') \times \pi_1(T^2)$ with $\mathbb{Z}^{n-1} \oplus \mathbb{Z}^2 = \mathbb{Z}^{n+1}$.

Let M be the graph manifold obtained by the gluings just described. It is readily seen that the natural projections $V_i \rightarrow \overline{N}_i$ define a projection $q: M \rightarrow D\overline{N}$, where $D\overline{N}$ is the double of the natural compactification of N . The map q is a locally trivial fiber-bundle with fibers homeomorphic to T^2 . If γ is the support of any simple curve joining the two boundary components of \overline{N} , then the double α of γ defines a simple loop in $D\overline{N}$. Let $L = q^{-1}(\alpha)$. It is easily seen that

$$\pi_1(L) \cong \langle x, y, z \mid yz = zy, xy = yzx, xz = zx \rangle \cong \mathbb{Z}^2 \rtimes_{\psi} \mathbb{Z},$$

where if x generates \mathbb{Z} we have $\psi(x)(y, z) = (y, y+z)$. Moreover, if L' is the intersection of L with one component Y of $\partial V_1 = \partial V_2 \subseteq M$, then $L' \cong T^2$, and $i: L' \rightarrow L$ induces an injective homomorphism $i_*: \pi_1(L') \rightarrow \pi_1(L)$ with $i_*(\pi_1(L')) = \langle y, z \rangle$. It is well-known (see *e.g.* [BrHa, III.Γ.4.17]) that i_* is *not* a quasi-isometric embedding, so the inclusion of $\pi_1(L')$ into $\pi_1(M)$ is *not* a quasi-isometric embedding (see Remark 1.1).

On the other hand, since the inclusion $\pi_1(L') \hookrightarrow \pi_1(Y)$ is a quasi-isometric embedding, if the inclusion $\pi_1(Y) \hookrightarrow \pi_1(M)$ were a quasi-isometric embedding, then by Remark 1.1 the inclusion $\pi_1(L') \hookrightarrow \pi_1(M)$ would also be quasi-isometric, while we have just proved that this is not the case. Therefore, the inclusion $\pi_1(Y) \hookrightarrow \pi_1(M)$ is also not a quasi-isometric embedding, and by the Milnor-Svarc Lemma, this implies that there exist walls of \widetilde{M} which are not quasi-isometrically embedded in \widetilde{M} .

As a consequence, M cannot support any locally CAT(0) metric: in fact, due to Milnor-Svarc Lemma and the Flat Torus Theorem (see *e.g.* [BrHa, pg. 475]), if a compact manifold M supports a locally CAT(0) metric and $H < \pi_1(M)$ is isomorphic to \mathbb{Z}^r for some $r \geq 1$, then H is necessarily quasi-isometrically embedded in $\pi_1(M)$. \square

We can exploit Proposition 2.19 to prove a portion of Theorem 0.15 in any dimension $n \geq 5$ (the case $n = 4$ will be dealt with in Chapter 12). Indeed, it is a well-known folk result that for every $n \geq 3$, there exists a cusped hyperbolic n -manifold with at least two cusps, and whose cusps are all toric. By taking finite covers, we can then obtain infinitely many such manifolds. McReynolds has informed the authors that a detailed proof of this folk result will be included in his forthcoming paper with Reid and Stover [MRS], where they will moreover show that in each dimension, such manifolds fall into infinitely many distinct commensurability classes. Applying Proposition 2.19, we immediately deduce:

COROLLARY 2.20. *For every $n \geq 5$, there exist infinitely many n -dimensional graph manifolds which do not support any locally CAT(0) metric.*

CHAPTER 3

Topological rigidity

In this chapter, we will establish various topological results for graph manifolds. The main goal will be to establish Theorem 0.2, which we restate here for the reader's convenience.

THEOREM (Borel Conjecture). *Let M be a graph manifold (possibly with boundary), of dimension $n \geq 6$. Assume M' is an arbitrary manifold and $\rho : M' \rightarrow M$ is a homotopy equivalence which restricts to a homeomorphism $\rho|_{\partial M'} : \partial M' \rightarrow \partial M$ between the boundaries of the manifolds. Then ρ is homotopic, rel ∂ , to a homeomorphism $\bar{\rho} : M' \rightarrow M$.*

In many of the arguments in this chapter, we will proceed by repeatedly cutting open a graph manifold along internal walls. Whenever we do such a cut, we effectively trade an internal wall in the original manifold for two boundary walls in the new manifold. Using the total number of internal walls as a measure of complexity, we see that this process *decreases* the overall complexity, allowing us to use inductive arguments in our proofs.

3.1. Contractible universal cover

A basic result in metric geometry implies that the universal cover of a closed CAT(0) manifold is contractible, and hence that any such manifold is aspherical. We establish the analogue:

LEMMA 3.1. *If M is a graph manifold (possibly with boundary), then M is aspherical.*

PROOF. We argue by induction on the number k of internal walls in M . If $k = 0$, then M is homeomorphic to the product $\bar{N} \times T^k$ where \bar{N} is a finite volume hyperbolic manifold with cusps cut off, and T^k is a torus. Since both factors are aspherical, and a product of aspherical manifolds is aspherical, the base case for our induction is established.

Now assume M has $k > 0$ internal walls, and that the result holds for all graph manifolds with $< k$ internal walls. Let W be an arbitrary internal wall in M , and cut M open along W . There are two cases to consider, according to whether W separates M into two components or not. We deal with the case where W separates M into M_1 and M_2 (the other case uses a similar reasoning). Since each of the M_i are graph manifolds, and have $< k$ internal walls, the inductive hypothesis ensures they are aspherical. So M is obtained by gluing together the two aspherical spaces M_1, M_2 along a common aspherical subspace W . A result of Whitehead [**Wh**] now asserts that M is also aspherical, *provided* that each of the inclusions $W \hookrightarrow M_i$ is π_1 -injective. But this is an easy consequence of the description of $\pi_1(M)$ as

the fundamental group of a graph of groups (see Section 2.2). This completes the inductive step, and establishes the Lemma. \square

Since the higher homotopy groups of our graph manifolds vanish, one would expect the topology of our manifolds to be determined by the fundamental group.

3.2. Lower algebraic K-theory

In the field of high-dimensional topology, some of the most important invariants of a manifold M are the (lower) algebraic K -groups of the integral group ring of the fundamental group. Obstructions to various natural problems often reside in these groups, and in some cases, all elements in the group can be realized as such obstructions. As a result, it is of some interest to obtain vanishing results for the lower K -groups. We will focus on the following covariant functors:

- the *Whitehead group* of M , $Wh(\pi_1(M))$, which is a quotient of the group $K_1(\mathbb{Z}[\pi_1(M)])$,
- the reduced K_0 -group, $\tilde{K}_0(\mathbb{Z}[\pi_1(M)])$, and
- the lower K -groups, $K_i(\mathbb{Z}[\pi_1(M)])$ with $i \leq -1$.

To simplify notation, we define the functors Wh_i (for $i \leq 1$) from the category of groups to the category of abelian groups as follows:

$$Wh_i(\Gamma) := \begin{cases} Wh(\Gamma) & i = 1 \\ \tilde{K}_0(\mathbb{Z}[\Gamma]) & i = 0 \\ K_i(\mathbb{Z}[\Gamma]) & i \leq -1 \end{cases}$$

When M is a closed manifold of non-positive sectional curvature of dimension $n \geq 5$, it follows from work of Farrell and Jones [**FaJo1**] that $Wh_i(\pi_1(M)) = 0$ for all $i \leq 1$. As a special case, $Wh_i(\mathbb{Z}^k)$ vanishes for $i \leq 1$, $k \geq 5$ (in fact, using work of Bass, Heller, and Swan [**BHS**] one can establish this for all k). We now generalize this vanishing result to graph manifolds.

PROPOSITION 3.2 (Lower K -groups vanish). *Let M be a graph manifold (possibly with boundary) of dimension $n \geq 6$ and $\Gamma = \pi_1(M)$. Then we have that $Wh_i(\Gamma) = 0$ for all $i \leq 1$.*

PROOF. We proceed by induction on the number k of internal walls in the graph manifold M . If $k = 0$, then M is a manifold with boundary, whose interior is homeomorphic to the product of a finite volume hyperbolic manifold with a torus. For such a manifold, Farrell and Jones [**FaJo2**] have already established the vanishing of Wh , \tilde{K}_0 , and K_i ($i \leq -1$).

So we may now assume that $k > 0$, and that the functors vanish for all graph manifolds with $< k$ internal walls. Pick an arbitrary internal wall in M , and consider the induced splitting of the group Γ . There are two cases to consider:

- (1) if the wall separates M into two components, then $\Gamma = G_1 *_H G_2$ is an amalgamation of two groups G_1, G_2 over a subgroup H .
- (2) if the wall does *not* separate, then $\Gamma = G *_H$ is isomorphic to an HNN extension of G over a subgroup H .

Moreover, $H \cong \mathbb{Z}^{n-1}$, and G, G_1, G_2 are fundamental groups of graph manifolds with $< k$ internal walls. We explain Case (1) in detail, as the argument for Case (2) is completely analogous.

Waldhausen has established [Wa3], [Wa4] (see also Bartels and Lück [BaLu] and Connolly and Prassidis [CoPr]) a Mayer-Vietoris type sequence for the functors Wh_i of an amalgamation $\Gamma = G_1 *_H G_2$ (or of an amalgamation $\Gamma = *_H G$). Waldhausen's sequence requires an "adjustment term" to $Wh_i(\Gamma)$, and takes the form:

$$(3.1) \quad \begin{aligned} \dots \rightarrow Wh_i(H) \rightarrow Wh_i(G_1) \oplus Wh_i(G_2) \rightarrow Wh_i(\Gamma)/Nil_i \\ \rightarrow Wh_{i-1}(H) \rightarrow Wh_{i-1}(G_1) \oplus Wh_{i-1}(G_2) \rightarrow \dots \end{aligned}$$

In the above sequence, the adjustment terms Nil_i are called the *Waldhausen Nil-groups* associated to the amalgamation $G_1 *_H G_2$.

For our specific amalgamation, the inductive hypothesis ensures that the terms involving the G_i all vanish. Since H is free abelian, the terms involving H also all vanish. Hence the Waldhausen long exact sequence gives us an isomorphism $Wh_i(\Gamma) \cong Nil_i$ for $i \leq -1$. Now the Waldhausen Nil-groups for a general amalgamation are extremely difficult to compute. However, when the amalgamating subgroup H has the property that its integral group ring $\mathbb{Z}[H]$ is *regular coherent*, Waldhausen has shown that the Nil-groups all vanish (see [Wa3, Theorem 4]). Recall that a ring R is said to be regular coherent provided every finitely generated R -module has a finite-length resolution by finitely generated projective R -modules. But it is an old result of Hall [Ha] that the integral group ring of finitely generated free abelian groups are regular coherent. This gives us $Wh_i(\Gamma) \cong Nil_i = 0$ for $i \leq 1$, concluding the inductive step in Case (1). In Case (2), we can apply an identical argument to the analogous long exact sequence for $\Gamma = G *_H$:

$$(3.2) \quad \begin{aligned} \dots \rightarrow Wh_i(H) \rightarrow Wh_i(G) \rightarrow Wh_i(\Gamma)/Nil_i \\ \rightarrow Wh_{i-1}(H) \rightarrow Wh_{i-1}(G) \rightarrow \dots \end{aligned}$$

This completes the proof of the proposition. \square

3.3. Borel Conjecture

Having established these preliminary results, we now turn to showing our main result for this section, Theorem 0.2. We are given a graph manifold M , of dimension ≥ 6 , and a homotopy equivalence $\rho : M' \rightarrow M$ where M' is an arbitrary manifold. Our goal is to find a homeomorphism $\bar{\rho} : M' \rightarrow M$ homotopic to ρ .

The proof of the theorem will again proceed by induction on k , the number of internal walls in the manifold M . The base case for our induction, $k = 0$, corresponds to a manifold with boundary, whose interior is homeomorphic to the product of a finite volume hyperbolic manifold with a torus. The Borel Conjecture for such manifolds follows from work of Farrell and Jones (see [FaJo2]). So we may now assume that $k > 0$, and choose an arbitrary internal wall inside M . Recall that this wall is a smoothly embedded codimension zero submanifold, homeomorphic to $T^{n-1} \times [0, 1]$. Let $N \subset M$ denote the smooth submanifold corresponding to the middle of the wall $T^{n-1} \times \{1/2\}$.

As a first step, we want to homotope the homotopy equivalence ρ to a smooth map $f : M' \rightarrow M$ having the property that (i) f is transverse to N , (ii) f restricts to a homotopy equivalence $f|_{f^{-1}(N)} : f^{-1}(N) \rightarrow N$, and (iii) f restricts to a homotopy equivalence from $M' \setminus f^{-1}(N)$ to $M \setminus N$. This question was studied by Cappell [Ca], who showed that there are two obstructions to being able to do this:

- an element in a suitable quotient group of $Wh(\pi_1(M))$, and

- an element in a group $UNil$ defined by Cappell, which depends on the decomposition of $\pi_1(M)$ as an amalgamation over $\pi_1(N)$ (or on the expression of $\pi_1(M)$ as an HNN-extension over $\pi_1(N)$).

Applying Proposition 3.2, we see that the first obstruction automatically vanishes in our situation. To deal with the second obstruction, we use a result of Cappell [Ca] showing that the $UNil$ group vanishes provided the subgroup $\pi_1(N)$ is *square-root closed* in the group $\pi_1(M)$. Recall that a subgroup $H \leq G$ is square-root closed provided that for $g \in G$, $g^2 \in H$ forces $g \in H$. So our next goal is to show that the fundamental group $\pi_1(N)$ of an internal wall is always square-root closed in the fundamental group $\pi_1(M)$ of the ambient graph manifold. This will follow from the:

Assertion: Let G be the fundamental group of a graph of groups. Assume that, in the graph of groups, each edge group is square-root closed in each adjacent vertex group. Then each edge group is square-root closed in G .

Assuming this **Assertion**, it is sufficient to show that the fundamental group of an internal wall $\pi_1(W)$ is square-root closed in the fundamental group of an adjacent chamber $\pi_1(C)$. From the product structure $C = \overline{N} \times T^k$ on the chambers, we have that $\pi_1(C)$ splits as a product $\pi_1(\overline{N}) \times \mathbb{Z}^k$, where \overline{N} is a suitable finite volume hyperbolic manifold with cusps cut off, and the \mathbb{Z}^k comes from the torus factor. W is a boundary component of C , hence splits as $\pi_1(Y) \times \mathbb{Z}^k$, where $Y \subset \overline{N}$ is a boundary component of \overline{N} . It is immediate from the definition that $\pi_1(W)$ is square-root closed in $\pi_1(C)$ if and only if $\pi_1(Y)$ is square-root closed in $\pi_1(\overline{N})$.

Using the induced action of $\pi_1(\overline{N})$ on the neutered space B (see Section 2.3), we can identify $\pi_1(Y)$ with the stabilizer of a boundary horosphere component \tilde{Y} in B . Now assume that $g \in \pi_1(\overline{N})$ satisfies $g^2 \in \pi_1(Y)$, but $g \notin \pi_1(Y)$. Then g^2 maps \tilde{Y} to itself, but g maps \tilde{Y} to some other boundary component $\tilde{Y}' \neq \tilde{Y}$, i.e. g interchanges the two horospheres \tilde{Y}' and \tilde{Y} . Since these two horospheres are centered at different points at infinity, there is a unique minimal length geodesic segment η joining \tilde{Y}' to \tilde{Y} . But g acts isometrically, and interchanges the two horospheres, hence must leave η invariant. This forces g to fix the midpoint of η , contradicting the fact that the $\pi_1(\overline{N})$ action on B is free. We conclude that every $\pi_1(W)$ is square-root closed in each adjacent $\pi_1(C)$.

So to obtain the vanishing of $UNil$, we are left with showing the **Assertion**. To this end, let T be the Bass-Serre tree corresponding to the graph of group, and consider the G -action on T . Choose an edge group $H \subset G$, which we can identify with the stabilizer of an edge $e \subset T$ in the tree T . Assume that $g \in G$ satisfies $g^2 \in H$ but $g \notin H$. This implies that $g \cdot e = e' \neq e$, but $g^2 \cdot e = e$. Again, we can let η be a minimal geodesic in the tree T joining the distinct edges e, e' . Since g interchanges the edges e, e' , it must leave invariant the geodesic η , hence fixes the midpoint of η . The action of G on the tree T is without inversions, so the midpoint of η must be a vertex v (i.e. η must be a path of even length). Let e_1, e_2 be the two edges in the path η adjacent to v (if the path is degenerate, take $e_1 = e$ and $e_2 = e'$). Since η is a geodesic, we have $e_1 \neq e_2$. But now we see that $g \in \text{Stab}(v)$ satisfies $g^2 \in \text{Stab}(e_1)$, but $g \notin \text{Stab}(e_1)$. This contradicts the assumption that each edge group is square-root closed in each adjacent vertex group, completing the proof of the **Assertion**.

From Cappell's work, we have now succeeded in homotoping the homotopy equivalence ρ to a map f with the property that f is transverse to the middle $T^{n-1} \times \{1/2\}$ of the thick wall $T^{n-1} \times [0, 1]$. Moreover, the homotopy can be chosen to have support away from the boundary $\partial M'$. Let $N' = f^{-1}(T^{n-1} \times \{1/2\})$, and note that by transversality N' is an $(n-1)$ -dimensional submanifold of M' . Denote by $F : N' \rightarrow T^{n-1}$ the restriction of f to N' . Since transversality is an open condition, we can assume (by a further perturbation if necessary) that f satisfies:

- f is transverse to *every* $T^{n-1} \times \{t\}$ for $t \in [0, 1]$,
- each $f^{-1}(T^{n-1} \times \{t\})$ is also homeomorphic to N' ,
- at the cost of reparametrizing, the restriction of f to the codimension zero submanifolds $N' \times [0, 1] \rightarrow T^{n-1} \times [0, 1]$ takes the form $F \times Id$.

The map $F : N' \rightarrow T^{n-1}$ is a homotopy equivalence, and as such, the manifold N' must be closed and connected (as these properties are detected by the homotopy invariant functors $H_{n-1}(-; \mathbb{Z}_2)$ and $\pi_0(-)$ respectively). The map F is a local diffeomorphism (by transversality of f) between a pair of closed smooth manifolds, hence is in fact a covering map. Finally, the degree of F must be $= 1$, since F induces an isomorphism on π_1 , so we conclude that F is a diffeomorphism between N' and T^{n-1} .

Now consider cutting M open along the submanifold $T^{n-1} \times \{1/2\}$. There are two possibilities, according to whether the complement of the wall has one or two connected components. We focus on the first case, since the second case is completely analogous. We now have a new manifold $M_0 := M \setminus (T^{n-1} \times \{1/2\})$ with two open ends, and we denote by \bar{M} the obvious compactification of M_0 obtained by closing off each end by attaching a copy of T^{n-1} . The compact manifold \bar{M} is a graph manifold, with one fewer internal wall than M , but with two additional boundary components. Likewise, we can cut M' open along the submanifold $N' \times \{1/2\}$, resulting in a manifold $M'_0 = M' \setminus (N' \times \{1/2\})$ with two open ends, and corresponding manifold with boundary \bar{M}' obtained from M'_0 by compactifying both ends with a copy of N' . Now the map f induces a map, which we denote g_0 , from M'_0 to M_0 . From the specific form of f in the vicinity of the submanifold $N' \subset M'$, we see that g_0 obviously extends to a map $g : \bar{M}' \rightarrow \bar{M}$ between the compactifications, which induces a homeomorphism (in fact a diffeomorphism) between the compactifying set $\bar{M}' \setminus M'_0$ (two copies of N^{n-1}) and the compactifying set $\bar{M} \setminus M_0$ (two copies of T^{n-1}). By Cappell's property (iii), g_0 is a homotopy equivalence, and since we have obvious homotopy equivalences $\bar{M}' \simeq M'_0$ and $\bar{M} \simeq M_0$, we conclude that g is also a homotopy equivalence.

We now have that \bar{M} is a graph manifold, with $< k$ internal walls, and a homotopy equivalence $g : \bar{M}' \rightarrow \bar{M}$ which restricts to a homeomorphism from $\partial \bar{M}'$ to $\partial \bar{M}$. From the inductive hypothesis, we see that the map g is homotopic, rel ∂ , to a homeomorphism. Since the homotopy leaves the boundaries unchanged, we can lift the homotopy, via the obvious "re-gluing" of boundary components, to a homotopy from $f : M' \rightarrow M$ to a new map $\bar{\rho} : M' \rightarrow M$. Moreover, it is immediate that the map $\bar{\rho}$ is a homeomorphism, completing the inductive step, and concluding the proof of our Theorem 0.2.

REMARK 3.3. The attentive reader will notice that the argument presented for Theorem 0.2 actually holds in considerably greater generality. The key steps of the proof are:

- establishing that the manifolds are aspherical, which only required the inclusion of each internal wall into each chamber to be π_1 -injective (in order to apply Whitehead's work),
- establishing the vanishing of the Whitehead group of the manifold, which required vanishing of the Whitehead group of the chambers, vanishing of the \tilde{K}_0 of the walls, and regular coherence of the fundamental group of the wall,
- vanishing of Cappell's $UNil$ obstruction, which follows from the square-root closed condition,
- knowledge of the Borel conjecture for each of the chambers.

All of these conditions can be locally verified, in that they are purely conditions on the individual chambers (and walls). By varying the chambers allowed, one can easily construct other types of manifolds for which the exact same argument establishes topological rigidity.

For a concrete example, Nguyen Phan [Ng] introduced the class of *cuspidally decomposable* manifolds. These manifolds are defined in a manner similar to our graph manifolds, but have pieces which are homeomorphic to finite volume negatively curved locally symmetric spaces with the cusps truncated. The walls are homeomorphic to infra-nil manifolds. It is easy to see that each of the above four conditions are satisfied for these pieces and walls, and hence the same argument establishes the Borel conjecture for the class of cuspidally decomposable manifolds.

As another example, one can extend the class of graph manifolds, by also allowing the presence of pieces which are of the form $\Sigma \times T^{n-2}$ (with Σ a non-compact, finite volume, hyperbolic surface). It is straightforward to check that each of the four conditions above also apply when we allow these new types of pieces, so again, the same argument establishes the Borel conjecture for this somewhat larger class of manifolds. Some further discussion of this broader class of manifolds can be found in the concluding Chapter 13.

3.4. Baum-Connes Conjecture and consequences

As our final result, we discuss the Baum-Connes conjecture for fundamental groups of graph manifolds. Recall that to any group G , one can associate its reduced group C^* -algebra $C_r^*(G)$ (see Section 8.2 for the definition). For a torsion-free group, the *Baum-Connes Conjecture* predicts that the complex K -homology of the classifying space BG coincides with the topological K -theory of $C_r^*(G)$. We will actually establish a somewhat stronger result known as the Baum-Connes conjecture with coefficients (the latter has certain useful inheritance properties). For a thorough discussion of this subject, we refer the reader to the book [MV] or the survey article [LuR].

PROPOSITION 3.4 (Baum-Connes conjecture). *For M a graph manifold (possibly with boundary), $\pi_1(M)$ satisfies the Baum-Connes conjecture (with coefficients).*

PROOF. The fundamental group of a finite volume hyperbolic manifold is a-T-menable (see [CCJJV] for an extensive discussion of this notion), and an extension of an a-T-menable group by an amenable group remains a-T-menable (see [CCJJV, Ex. 6.1.6]). This tells us that the fundamental groups of pieces are always a-T-menable. But groups which are a-T-menable satisfy the Baum-Connes conjecture with coefficients (see Higson and Kasparov [HK, Thm. 1.1]). So $\pi_1(M)$ is the

fundamental group of a graph of groups, where all the vertex groups in the graph of groups satisfy the Baum-Connes conjecture with coefficients. This implies that $\pi_1(M)$ itself satisfies the Baum-Connes conjecture with coefficients (by work of Oyono-Oyono [O-O, Thm. 1.1]). Finally, the conjecture with coefficients implies the original Baum-Connes conjecture, completing the proof. \square

A nice feature of the Baum-Connes conjecture is that it is known to imply several other well-known conjectures. We explicitly mention three of these consequences which may be of general interest. Throughout the rest of this section, we let G denote the fundamental group of an arbitrary graph manifold.

COROLLARY 3.5 (Idempotent conjectures). *The Kadison Conjecture holds: the reduced C^* -algebra $C_r^*(G)$ has no idempotents. As a consequence, the Kaplansky Conjecture also holds: the group algebra $\mathbb{Q}G$ has no idempotents.*

COROLLARY 3.6 (Gromov-Lawson-Rosenberg conjecture). *Let W be a closed, connected, smooth, Spin manifold with $\pi_1(W) \cong G$. If W supports a Riemannian metric of positive scalar curvature, then the higher \hat{A} -genera of W all vanish.*

COROLLARY 3.7 (Zero-in-the-Spectrum conjecture). *Let M be a graph manifold, equipped with an arbitrary Riemannian metric. Then there exists some $p \geq 0$ so that zero lies in the spectrum of the Laplace-Beltrami operator Δ_p acting on square-integrable complex valued p -forms on \tilde{M} (the universal cover of M).*

Isomorphisms preserve pieces

This chapter is devoted to the proof of Theorem 0.3. We recall the statement here for convenience.

THEOREM. *Let M_1, M_2 be a pair of graph manifolds and let $\Gamma_i = \pi_1(M_i)$ be their respective fundamental groups. Let $\Lambda_1 \leq \Gamma_1$ be a subgroup conjugate to the fundamental group of piece V in M_1 , and $\phi : \Gamma_1 \rightarrow \Gamma_2$ be an isomorphism. Then $\phi(\Lambda_1)$ is conjugate to the fundamental group $\Lambda_2 \leq \Gamma_2$ of a piece in M_2 .*

Roughly speaking, the strategy is as follows. First, via the study of the action of Γ_i on its Bass-Serre tree, we show that ϕ has to take the fundamental group of a boundary component of a piece of M_1 onto (the conjugate of) the fundamental group of a boundary component of a piece of M_2 . This step requires that ϕ is an isomorphism, rather than simply a quasi-isometry.

We then use this result to show that the fundamental group Λ_1 of a piece of M_1 is taken by ϕ onto a subset of Γ_2 which lies at finite Hausdorff distance from (the conjugate of) the fundamental group Λ_2 of a piece of M_2 . Finally, a standard trick (using again the fact that ϕ is an isomorphism) allows us to conclude that this subset indeed coincides with a conjugate of Λ_2 .

4.1. Edge stabilizers are preserved by isomorphisms

Let M be a n -dimensional graph manifold with fundamental group Γ , and let T be the Bass-Serre tree of Γ associated to the realization of \widehat{M} as a tree of spaces. The following lemma provides an intrinsic group-theoretic characterization of the stabilizers of the edges of T .

LEMMA 4.1. *Let H be a subgroup of Γ isomorphic to \mathbb{Z}^{n-1} . Then H is contained in the stabilizer of an edge of T .*

PROOF. We first show that it is sufficient to prove that H is contained in the stabilizer of a vertex. In fact, if this is the case, let us denote by N and T^k respectively the hyperbolic and the toric factor of the piece of M corresponding to the vertex fixed by H . Notice that $k \leq n - 3$. We have that H is contained in a conjugate of $\pi_1(N \times T^k) < \pi_1(M)$. For our purposes, we can safely assume $H < \pi_1(N \times T^k)$. The projection of H on $\pi_1(N)$ is an abelian group of rank at least $n - k - 1 \geq 2$, and it is therefore contained in a cusp subgroup. This implies the conclusion.

Therefore, we only need to prove that H is contained in the stabilizer of a vertex. As H is a finitely generated nilpotent group, a standard result about groups acting on a tree (see [Se, Proposition 6.5.27]), guarantees that if H does *not* stabilize a vertex, then there exists a geodesic γ in T that is invariant under the action of H . So we only need to prove that there is no such geodesic.

If there was such a geodesic, we could find a subgroup $K < H$ of rank $n-2$ which fixes γ pointwise. Let v be any vertex of γ and let N and T^k be the hyperbolic and toric factor of the corresponding piece. As before, let us just assume $K < \pi_1(N \times T^k)$. Let $j: \pi_1(N \times T^k) \rightarrow \mathbb{Z}^k$ be the projection on the fiber subgroup. As the rank of K is $n-2$ and $n-2 > k$, the restriction $j|_K$ cannot be injective. That is to say, the group $K \cap \ker j$ contains a non-trivial element g . The element g acts as a deck transformation on the universal covering $B \times \mathbb{R}^k$ of $\overline{N} \times T^k$, where B is the neutered space covering \overline{N} . Since g belongs to $\ker j$, the action of g is trivial on the Euclidean factor of $B \times \mathbb{R}^k$. Moreover, since g belongs to K the action of g on B fixes two boundary components of B . As a consequence, g must fix a point in $B \times \mathbb{R}^k$ (see the proof of the **Assertion** in Section 3.3). This is a contradiction, as non-trivial deck transformations do not have fixed points. \square

Let us come back to the notation of the statement of Theorem 0.3. As a consequence of the previous lemma we readily get the following:

COROLLARY 4.2. *The isomorphism φ has to take the fundamental group H_1 of a boundary component of a piece of M_1 onto (the conjugate of) the fundamental group of a boundary component of a piece of M_2 .*

PROOF. Since H_1 is isomorphic to \mathbb{Z}^{n-1} , the previous lemma implies that $\varphi(H_1)$ is contained in the stabilizer of an edge of the Bass-Serre tree of Γ_2 , *i.e.* it is contained in gH_2g^{-1} , where H_2 is the fundamental group of a boundary component of a piece of M_2 and g is an element of Γ_2 . In order to conclude we have to show that $\varphi(H_1)$ actually coincides with gH_2g^{-1} . However, the subgroup $\varphi^{-1}(gH_2g^{-1})$ of Γ_1 obviously contains H_1 , and is contained in the stabilizer of an edge of the Bass-Serre tree of Γ_1 by Lemma 4.1 again. This forces $\varphi^{-1}(gH_2g^{-1}) = H_1$, whence the conclusion. \square

4.2. Isomorphisms quasi-preserve pieces

Let us consider the universal coverings $\widetilde{M}_1, \widetilde{M}_2$ of M_1, M_2 . By Milnor-Svarc's Lemma, φ induces a (k, c) -quasi-isometry $f: \widetilde{M}_1 \rightarrow \widetilde{M}_2$. By Corollary 4.2, together with the fact that the $\pi_1(M_i)$ -orbits of the walls of \widetilde{M}_i are in finite number, there exists a constant $\lambda > 0$ such that for every wall $W_1 \subseteq \widetilde{M}_1$ the set $f(W_1)$ is at Hausdorff distance bounded by λ from a wall $W_2 \subseteq \widetilde{M}_2$ (the wall W_2 is unique in view of Lemma 2.17).

The following result plays an important role in the proof both of Theorem 0.3 and of Theorem 0.12.

PROPOSITION 4.3. *Let $f: \widetilde{M}_1 \rightarrow \widetilde{M}_2$ be a (k, c) -quasi-isometry. Suppose there exists λ with the property that, for each wall W_1 of \widetilde{M}_1 , there exists a wall W_2 of \widetilde{M}_2 with the Hausdorff distance between $f(W_1)$ and W_2 bounded by λ .*

Then there exists a universal constant H with the property that, for every chamber $C_1 \subseteq \widetilde{M}_1$, there exists a unique chamber $C_2 \subseteq \widetilde{M}_2$ with the Hausdorff distance between $f(C_1)$ and C_2 bounded by H . Moreover, if W_1 is a wall adjacent to C_1 then $f(W_1)$ lies at finite Hausdorff distance from a wall W_2 adjacent to C_2 .

PROOF. Let us fix a chamber C_1 of \widetilde{M}_1 , and let W_1, W'_1 be walls adjacent to C_1 . Our hypothesis implies that there exist walls W_2, W'_2 of \widetilde{M}_2 such that $f(W_1)$ and $f(W'_1)$ lie within finite Hausdorff distance respectively from W_2 and W'_2 (such

walls are uniquely determined – see Lemma 2.17). We first prove that a chamber C_2 exists such that W_2 and W'_2 are both adjacent to C_2 .

Suppose by contradiction that there exists a wall $P_2 \subseteq \widetilde{M}_2$ such that $P_2 \neq W_2, W'_2$, and every continuous path connecting W_2, W'_2 intersects P_2 . Since f is a quasi-isometry, there exists a wall $P_1 \subseteq \widetilde{M}_1$ such that $f(P_1)$ is at Hausdorff distance at most λ from P_2 (just take P_1 to be the wall at bounded distance from $g(P_2)$, where g is a quasi-inverse to f).

We claim that for every $x > 0$, there exist points $w_1 \in W_1$, $w'_1 \in W'_1$ which are joined by a path $\gamma_1 : [0, l] \rightarrow \widetilde{M}_1$, parametrized by arc length, such that $d(\gamma_1(t), P_1) \geq x$ for every $t \in [0, l]$. First, observe that $P_1 \neq W_1, W'_1$. The realization of \widetilde{M} as a tree of spaces makes it clear that we might as well assume that P_1 is adjacent to C_1 . In this case, if $\pi : C_1 \rightarrow B_1$ is the projection of C_1 on its base, then $\pi(W_1) = O$, $\pi(W'_1) = O'$ and $\pi(P_1) = O''$ for distinct horospheres O, O', O'' of the neutered space $B_1 \subseteq \mathbb{H}^k$. Let us fix an identification of \mathbb{H}^k with the half-space model, in such a way that O'' corresponds to a horosphere centered at the point at infinity. Since $k \geq 3$, it is now easy to show that for every sufficiently small $\varepsilon > 0$ it is possible to join a point in O with a point in O' by a rectifiable path supported on the intersection of B_1 with the Euclidean horizontal hyperplane at height ε . In fact, this intersection is (homeomorphic to) \mathbb{R}^{k-1} with a countable family of open disjoint balls removed (recall that $k - 1 \geq 2$). Let $\gamma_1 : [0, l] \rightarrow C_1 \subseteq \widetilde{M}_1$ be the parameterization by arc length (with respect to the global distance on \widetilde{M}) of a lift to C_1 of such a path. It is clear that $d_{C_1}(\gamma_1(t), P_1) \geq x(\varepsilon)$ for every $t \in [0, l]$, where $x(\varepsilon)$ tends to $+\infty$ as ε tends to 0. Using Lemma 2.16 we obtain that the same holds for $d(\gamma_1(t), P_1)$, up to changing the function $x(\varepsilon)$.

Let γ_2 be the continuous path obtained by concatenating the geodesic segments joining $f(\gamma_1(n))$ to $f(\gamma_1(n+1))$, $0 \leq n \leq l$. The starting point of the obtained path is at a distance at most λ from W_2 and its final point is at a distance at most $\lambda + k + c$ from W'_2 . Moreover, for every t we have the estimate:

$$d(\gamma_2(t), P_2) \geq \frac{x(\varepsilon) - \lambda}{k} - k - 2c.$$

If $x(\varepsilon)$ is chosen big enough, we obtain $d(\gamma_2(t), P_2) \geq \lambda + k + c + 1$ for every t . By concatenating γ_2 with two geodesics joining its initial point with W_2 and its endpoint with W'_2 we obtain a continuous path connecting W_2, W'_2 which does not intersect P_2 : a contradiction. We have thus proved that every wall adjacent to C_1 is taken by f onto a subset which lies at Hausdorff distance bounded by λ from a wall adjacent to C_2 .

Let us now prove that C_2 lies at a universally bounded Hausdorff distance from $f(C_1)$. Since walls are h -dense in \widetilde{M} for some $h > 0$, for every $p_1 \in C_1$ there exists $p'_1 \in W_1$ with $d(p_1, p'_1) \leq h$, where W_1 is a wall adjacent to C_1 . Then

$$d(f(p_1), C_2) \leq d(f(p_1), f(p'_1)) + d(f(p'_1), C_2) \leq kh + c + \lambda.$$

This tells us that $f(C_1)$ is contained in the $(kh + c + \lambda)$ -neighbourhood of C_2 . Let g be the quasi-inverse of f . The same argument shows that $g(C_2)$ is contained in the $(kh + c + \lambda)$ -neighbourhood of some chamber C'_1 , and Lemma 2.18 implies that $C'_1 = C_1$. Now, if $q_2 \in C_2$ we have $d(q_2, f(g(q_2))) \leq c$, and there exists $q_1 \in C_1$

with $d(g(q_2), q_1) \leq kh + c + \lambda$. We now can estimate the distance

$$\begin{aligned} d(q_2, f(q_1)) &\leq d(q_2, f(g(q_2))) + d(f(g(q_2)), f(q_1)) \\ &\leq c + kd(g(q_2), q_1) + c \\ &\leq 2c + k(kh + c + \lambda). \end{aligned}$$

So we can set $H = k^2h + (k + 2)c + k\lambda$, and we are done. Finally, the uniqueness of C_2 is a consequence of Lemma 2.18. \square

4.3. Isomorphisms preserve pieces

We are ready to establish Theorem 0.3.

PROOF. By Proposition 4.3 and the Milnor-Svarc Lemma, if $\Lambda_1 < \Gamma_1$ is the fundamental group of a piece, then the Hausdorff distance between $\varphi(\Lambda_1)$ and $g\Lambda_2g^{-1}$ is bounded by H for some fundamental group of a piece $\Lambda_2 < \Gamma_2$ and some $g \in \Gamma_2$. Up to conjugation, and increasing H by $d(g, id)$, we may assume $g = id$.

A standard argument now allows us to prove that $\varphi(\Lambda_1) = \Lambda_2$. In fact, if $h \in \Lambda_1$ we have that

$$\varphi(h) \cdot \varphi(\Lambda_1) = \varphi(h \cdot \Lambda_1) = \varphi(\Lambda_1).$$

Since $\varphi(\Lambda_1)$ is at bounded Hausdorff distance from Λ_2 , this implies that $\varphi(h) \cdot \Lambda_2$ is at bounded Hausdorff distance from Λ_2 . By Milnor-Svarc's Lemma, if C_2 is the chamber of \widetilde{M}_2 that is fixed by Λ_2 , then the chamber $\varphi(h)(C_2)$ is at finite Hausdorff distance from C_2 . By Lemma 2.18 this implies in turn that $\varphi(h)(C_2) = C_2$, so $\varphi(h) \in \Lambda_2$, and $\varphi(\Lambda_1) \subseteq \Lambda_2$. Finally, since φ^{-1} is a quasi-inverse of φ , we have that $\varphi^{-1}(\Lambda_2)$ stays at finite distance from Λ_1 . The above argument again shows that $\varphi^{-1}(\Lambda_2) \subseteq \Lambda_1$. We conclude that $\varphi(\Lambda_1) = \Lambda_2$, completing the proof of Theorem 0.3. \square

Putting together Theorem 0.3 and Lemma 2.14 one can easily refine the statement of Theorem 0.3 as follows:

THEOREM 4.4. *Let M, M' be a pair of graph manifolds which decompose into pieces V_1, \dots, V_h , and V'_1, \dots, V'_k respectively. Suppose that $\varphi: \pi_1(M) \rightarrow \pi_1(M')$ is an isomorphism. Then $h = k$ and, up to reordering the indices, for every $i = 1, \dots, h$ the image of $\pi_1(V_i)$ under φ coincides with a conjugate of $\pi_1(V'_i)$. Moreover, with this choice of indices V_i is adjacent to V_j if and only if V'_i is adjacent to V'_j .*

As this will be useful in some later chapters, we point out the following easy generalization of Theorem 0.3:

THEOREM 4.5. *Let M_1, M_2 be a pair of graph manifolds and let $\Gamma_i = \pi_1(M_i)$ be their respective fundamental groups. Let $\Lambda_1 \leq \Gamma_1$ be a subgroup conjugate to the fundamental group of piece V in M_1 , and $\varphi: \Gamma_1 \rightarrow \Gamma_2$ be an injective homomorphism whose image has finite index in Γ_2 . Then $\varphi(\Lambda_1)$ is contained in the conjugate of the fundamental group of a piece in M_2 .*

PROOF. Being injective and having image of finite index, the homomorphism φ defines a quasi-isometry between Γ_1 and Γ_2 . Applying the Milnor-Svarc Lemma, φ induces a quasi-isometry f between \widetilde{M}_1 and \widetilde{M}_2 .

Lemma 4.1 implies that φ takes the fundamental group H_1 of any boundary component of a piece of M_1 into the conjugate H_2 of the fundamental group of a boundary component of a piece of M_2 . Also observe that the rank of the free abelian group H_1 coincides with the rank of H_2 , so the subgroup $\varphi(H_1)$ has finite index in H_2 . This implies that f satisfies the hypotheses of Proposition 4.3, which in turn implies that f quasi-preserved walls and chambers. Finally, the argument in the proof of Theorem 0.3 implies that φ takes the stabilizer of a chamber of \widetilde{M}_1 into the stabilizer of a chamber of \widetilde{M}_2 , as desired. \square

Smooth rigidity

This chapter is devoted to the proof of Theorem 0.5, which we recall here for the convenience of the reader:

THEOREM. *Let M, M' be graph manifolds and let $\varphi: \pi_1(M) \rightarrow \pi_1(M')$ be a group isomorphism. Then φ is induced by a diffeomorphism $\psi: M \rightarrow M'$.*

It will be clear from our construction that the diffeomorphism ψ of the above theorem can be chosen in such a way that $\psi|_{\partial M}: \partial M \rightarrow \partial M'$ is an affine diffeomorphism.

Since we will need to be careful about some well-known, but somewhat subtle, details of the theory of fundamental groups, we recall here some basic facts. If $f: M \rightarrow N$ is a continuous map between path connected spaces, then f induces a homomorphism $f_*: \pi_1(M) \rightarrow \pi_1(N)$ which is well-defined up to conjugacy (in $\pi_1(N)$). This is due to the fact that, for $x_0, x_1 \in M$, $x_0 \neq x_1$, the identification of $\pi_1(M, x_0)$ with $\pi_1(M, x_1)$ is canonical up to conjugacy, and the same holds when choosing different basepoints in N . If $\varphi: \pi_1(M) \rightarrow \pi_1(N)$ is a homomorphism, we will say that φ is induced by f if for some (and hence every) choice of basepoints $x_0 \in M$, $y_0 \in N$ the homomorphism $f_*: \pi_1(M, x_0) \rightarrow \pi_1(N, y_0)$ is equal to φ , up to conjugacy by an element of $\pi_1(N)$ (by the discussion above, this notion is indeed well-defined). Also observe that if V is a path connected subset of M and $i: V \hookrightarrow M$ is the inclusion, then we can define $i_*(\pi_1(V))$ as a subgroup of $\pi_1(M)$, well-defined up to conjugacy. When saying that $\pi_1(V)$ is a subgroup of $\pi_1(M)$, we will be implicitly choosing a preferred representative among the conjugate subgroups representing the conjugacy class of $\pi_1(V)$: this amounts to choosing a basepoint in V , a basepoint in M and a path joining these basepoints.

5.1. Defining ψ on the pieces

Let M, M' be a pair of graph manifolds, and suppose that $\varphi: \pi_1(M) \rightarrow \pi_1(M')$ is an isomorphism. By Theorem 4.4, the isomorphism φ establishes a bijection between the (conjugacy classes of the) fundamental groups of the pieces of M and M' . Let N_1, \dots, N_h (resp. N'_1, \dots, N'_h) be the (truncated) cusped hyperbolic manifolds such that $V_i = N_i \times T^{a_i}$ (resp. $V'_i = N'_i \times T^{b_i}$) are the pieces of M (resp. of M'), $i = 1, \dots, h$. From now on, for every $i = 1, \dots, h$, we fix an identification of $\pi_1(V_i)$ (resp. of $\pi_1(V'_i)$) with a distinguished subgroup of $\pi_1(M)$ (resp. of $\pi_1(M')$). As mentioned above, such an identification depends on the choice of one basepoint for M, M' and for each piece, and suitable paths connecting the basepoint of the whole manifolds with the basepoints of their pieces. We also fix $g_i \in \pi_1(M')$ such that $\varphi(\pi_1(V_i)) = g_i \pi_1(V'_i) g_i^{-1}$ for every $i = 1, \dots, h$.

LEMMA 5.1. *For $i = 1, \dots, h$ there exists a diffeomorphism $\psi_i: V_i \rightarrow V'_i$ which induces the isomorphism $g \mapsto g_i^{-1}\varphi(g)g_i$ between $\pi_1(V_i)$ and $\pi_1(V'_i)$, and restricts to an affine diffeomorphism of ∂V_i onto $\partial V'_i$.*

PROOF. Set $V = V_i$, $V' = V'_i$, $N = N_i$, $N' = N'_i$. The center of $\pi_1(V)$ is equal to the fundamental group of its toric factor (see Remark 2.9), so $\pi_1(N)$ is just the quotient of $\pi_1(V)$ by its center, and the same holds true for $\pi_1(N')$. We have in particular $V = N \times T^a$, $V' = N' \times T^a$ for the same $a \in \mathbb{N}$, so $\pi_1(V)$ (resp. $\pi_1(V')$) is canonically isomorphic to $\pi_1(N) \times \mathbb{Z}^a$ (resp. $\pi_1(N') \times \mathbb{Z}^a$), and the isomorphism $\varphi_i: \pi_1(V) \rightarrow \pi_1(V')$ defined by $\varphi_i(g) = g_i^{-1}\varphi(g)g_i$ for every $g \in \pi_1(V)$ induces an isomorphism $\theta: \pi_1(N) \rightarrow \pi_1(N')$. Henceforth, we identify T^a with the quotient of \mathbb{R}^a by the standard action of \mathbb{Z}^a , *i.e.* we fix an identification of $\pi_1(T^a)$ with $\mathbb{Z}^a \subseteq \mathbb{R}^a$ (since $\pi_1(T^a)$ is abelian, we do not need to worry about choice of basepoints). Then the isomorphism $\varphi_i: \pi_1(N) \times \mathbb{Z}^a \rightarrow \pi_1(N') \times \mathbb{Z}^a$ has the form $\varphi_i(g, v) = (\theta(g), f(g, v))$, where $f: \pi_1(N) \times \mathbb{Z}^a \rightarrow \mathbb{Z}^a$ is a homomorphism. If $\beta: \mathbb{Z}^a \rightarrow \mathbb{Z}^a$, $\alpha: \pi_1(N) \rightarrow \mathbb{Z}^a$ are defined by $\beta(v) = f(1, v)$ and $\alpha(g) = f(g, 0)$, we have that

$$\varphi_i(g, v) = (\theta(g), \alpha(g) + \beta(v)) \quad \text{for every } g \in \pi_1(N), v \in \mathbb{Z}^a.$$

Moreover, since φ_i is an isomorphism, we have that α is a homomorphism and β is an automorphism. Any automorphism of $\pi_1(T^a)$ is induced by an affine diffeomorphism of T^a onto itself, so in order to construct the required diffeomorphism $\psi: V \rightarrow V'$ inducing φ_i it is not restrictive to assume that $\beta(v) = v$ for every $v \in \mathbb{Z}^a$.

Let us now fix identifications $\pi_1(N) \cong \Gamma < \text{Isom}(\mathbb{H}^l)$, $\pi_1(N') \cong \Gamma' < \text{Isom}(\mathbb{H}^l)$, $N = B/\Gamma$, $N' = B'/\Gamma'$, where $B, B' \subseteq \mathbb{H}^l$ are the neutered spaces providing the universal coverings of N, N' . For later purposes, we will denote by $p_1: B \rightarrow N$ the covering map just introduced. Mostow rigidity provides an isometry $\tilde{\kappa}: \mathbb{H}^l \rightarrow \mathbb{H}^l$ such that $\tilde{\kappa}(g \cdot x) = \theta(g) \cdot \tilde{\kappa}(x)$ for every $g \in \Gamma$, $x \in \mathbb{H}^l$. Up to changing the choice of the horospherical sections defining N as the truncation of a cusped hyperbolic manifold, we may also suppose that $\tilde{\kappa}(B) = B'$ (see Remark 2.3). We now establish the following:

Claim: There exists a smooth function $\tilde{\eta}: B \rightarrow \mathbb{R}^a$ such that $\tilde{\eta}(g \cdot x) = \tilde{\eta}(x) + \alpha(g)$ for every $x \in B$, $g \in \Gamma$.

In fact, let Γ act on $B \times \mathbb{R}^a$ by setting $g \cdot (x, v) = (g \cdot x, v + \alpha(g))$, denote by Y the quotient space and let $p_2: B \times \mathbb{R}^a \rightarrow Y$ be the natural projection. Since N is canonically identified with the quotient of B by the action of Γ , we have a canonical projection $p_3: Y \rightarrow N$, which defines a natural structure of flat affine fiber bundle. More precisely, Y is the total space of a flat fiber bundle with fiber \mathbb{R}^a and structural group given by the group of integer translations of \mathbb{R}^a . In particular, every fiber of p_3 inherits a well-defined affine structure, so it is possible to define affine combinations of points in a fiber. Exploiting this fact, we can use a suitable partition of unity to glue local sections of p_3 into a global smooth section $s: N \rightarrow Y$.

$$\begin{array}{ccc} B \times \mathbb{R}^a & \xrightarrow{p_2} & Y \\ & & \downarrow p_3 \\ B & \xrightarrow{p_1} & N \end{array} \quad \left. \vphantom{\begin{array}{ccc} B \times \mathbb{R}^a & \xrightarrow{p_2} & Y \\ & & \downarrow p_3 \\ B & \xrightarrow{p_1} & N \end{array}} \right\} s$$

We now define η as follows. Let us take $x \in B$. For every $v \in \mathbb{R}^a$ we have $p_3(p_2(x, v)) = p_1(x)$. Moreover, by construction $p_2(x, v) = p_2(x, w)$ if and only if $v = w$. As a consequence, there exists a unique $\tilde{\eta}(x) \in \mathbb{R}^a$ such that $p_2(x, \tilde{\eta}(x)) = s(p_1(x))$. Since p_1, p_2, s are smooth, $\tilde{\eta}$ is also smooth. Moreover, for $x \in B$ and $g \in \Gamma$ we have:

$$\begin{aligned} p_2(g \cdot x, \tilde{\eta}(x) + \alpha(g)) &= p_2(g \cdot (x, \tilde{\eta}(x))) = p_2(x, \tilde{\eta}(x)) \\ &= s(p_1(x)) = s(p_1(g \cdot x)). \end{aligned}$$

The first equality is due to the definition of the Γ -action on $B \times \mathbb{R}^a$. The second and fourth equality are immediate from the definition of the quotient maps p_2 and p_1 respectively. The third equality follows from the choice of $\tilde{\eta}$ (see previous paragraph). Finally, comparing the first and last term, we see that $\tilde{\eta}(x) + \alpha(g)$ satisfies the defining property for the point $\tilde{\eta}(g \cdot x)$, so by uniqueness we obtain $\tilde{\eta}(g \cdot x) = \tilde{\eta}(x) + \alpha(g)$, and the **Claim** is proved.

We now return to the proof of the Lemma. Define the map $\tilde{\psi}: B \times \mathbb{R}^a \rightarrow B' \times \mathbb{R}^a$ via $\tilde{\psi}(x, v) = (\tilde{\kappa}(x), v + \tilde{\eta}(x))$. Of course $\tilde{\psi}$ is a diffeomorphism. Moreover, for every $(x, v) \in B \times \mathbb{R}^a$ and $(g, w) \in \Gamma \times \mathbb{Z}^a \cong \pi_1(V)$, we have

$$\begin{aligned} \tilde{\psi}((g, w) \cdot (x, v)) &= \tilde{\psi}(g \cdot x, v + w) \\ &= (\tilde{\kappa}(g \cdot x), v + w + \tilde{\eta}(g \cdot x)) \\ &= (\theta(g) \cdot \tilde{\kappa}(x), v + w + \tilde{\eta}(x) + \alpha(g)) \\ &= (\theta(g), w + \alpha(g)) \cdot (\tilde{\kappa}(x), v + \tilde{\eta}(x)) \\ &= (\theta(g), w + \alpha(g)) \cdot \tilde{\psi}(x, v) \end{aligned}$$

so $\tilde{\psi}$ induces a diffeomorphism $\bar{\psi}: V \rightarrow V'$ inducing the isomorphism φ at the level of fundamental groups. Now let $\kappa: N \rightarrow N'$ be the isometry induced by $\tilde{\kappa}$, H be a component of ∂N , and set $H' = \kappa(H) \subseteq \partial N'$. By construction, the restriction of $\bar{\psi}$ to the component $H \times T^a$ of ∂V has the form

$$H \times T^a \rightarrow H' \times T^a, \quad (x, v) \mapsto (\kappa(x), v + \bar{\eta}(x))$$

for some smooth $\bar{\eta}: H \rightarrow T^a$. Recall that H is affinely diffeomorphic to a torus T^b , and that every map between affine tori is homotopic to an affine map, so $\bar{\eta}$ is homotopic to an affine map $\eta: H \rightarrow T^a$. Using this homotopy, we modify $\bar{\psi}$ in a collar of $H \times T^a$ in order to get a diffeomorphism $\psi: V \rightarrow V'$ whose restriction to $H \times T^a$ has the form $(x, v) \mapsto (\kappa(x), v + \eta(x))$. After repeating this procedure for every component of ∂V we are left with the desired diffeomorphism ψ . \square

5.2. Dehn twists

We now would like to glue the diffeomorphisms $\psi_i: V_i \rightarrow V'_i$ provided by Lemma 5.1 into a diffeomorphism $\psi: M \rightarrow M'$. In order to make this strategy work, we have to take care of two issues. First, to define ψ we have to check that if V_i and V_j share a boundary component H , then ψ_i and ψ_j coincide on H . Once this has been established, we have to ensure that the obtained ψ induces the isomorphism $\varphi: \pi_1(M) \rightarrow \pi_1(M')$ fixed at the beginning of the section. The following

remark, which is essentially due - in a different context - to Nguyen Phan [Ng], shows that the issues just discussed may really hide some subtleties.

REMARK 5.2. Suppose $M = M'$ is a graph manifold obtained by gluing two pieces V_1, V_2 along their unique boundary component $H = V_1 \cap V_2 \subseteq M$. Fix a basepoint $x_0 \in H$, and set $G_1 = \pi_1(V_1, x_0)$, $G_2 = \pi_1(V_2, x_0)$, $K = \pi_1(H, x_0)$. The group $\pi_1(M, x_0)$ is canonically identified with the amalgamated product $G = G_1 *_K G_2$, where we are considering K as a subgroup of G_1 and G_2 via the natural (injective) maps induced by the inclusions $H \hookrightarrow V_1$, $H \hookrightarrow V_2$. Let us take $g_0 \in K \setminus \{1\}$. Since K is abelian, there exists a unique isomorphism $\varphi: G \rightarrow G$ such that

$$\varphi(g) = \begin{cases} g & \text{if } g \in G_1 \\ g_0 g g_0^{-1} & \text{if } g \in G_2 \end{cases}.$$

It is easy to see that, in this special case, the construction described in Lemma 5.1 leads to diffeomorphisms $\psi_1: V_1 \rightarrow V_1$, $\psi_2: V_2 \rightarrow V_2$ which can be chosen to equal the identity on V_1, V_2 respectively. In particular, since M and M' are obtained by gluing V_1 and V_2 exactly in the same way, no issue about the possibility of defining ψ arises. However, if we chose naively to glue ψ_1 and ψ_2 simply by requiring that $\psi|_{V_i} = \psi_i$, we would obtain $\psi = \text{Id}_M$. But this contradicts the fact that, when the element g_0 is chosen appropriately, φ may define a non-trivial outer automorphism of G (of infinite order), see Lemma 5.4 below.

The previous remark motivates the following:

DEFINITION 5.3. Suppose M is a graph manifold, and let V_1, V_2 be pieces of M glued to each other along a common toric component H of ∂V_1 and ∂V_2 . Let h be a fixed element of $\pi_1(H)$ (since $\pi_1(H)$ is abelian, this is independent of basepoints). The *Dehn twist t_h along h* is the diffeomorphism $t_h: M \rightarrow M$ which is defined as follows.

By construction, H admits a collar U in M which is canonically foliated by tori (see Chapter 2). In particular, U is affinely diffeomorphic to $T^{n-1} \times [-1, 1]$, where $T^{n-1} = \mathbb{R}^{n-1}/\mathbb{Z}^{n-1}$ is the standard affine $(n-1)$ -torus, and $\pi_1(H)$ is canonically identified with the group \mathbb{Z}^{n-1} of the automorphisms of the covering $\pi: \mathbb{R}^{n-1} \rightarrow T^{n-1}$. Let now $l: [-1, 1] \rightarrow [0, 1]$ be a smooth function such that $l|_{[-1, -1+\varepsilon]} = 0$, $l|_{[1-\varepsilon, 1]} = 1$ and set

$$\tilde{t}_h: \mathbb{R}^{n-1} \times [-1, 1] \rightarrow \mathbb{R}^{n-1} \times [-1, 1], \quad \tilde{t}_h(v, s) = (v + l(s) \cdot h, s).$$

The map \tilde{t}_h is \mathbb{Z}^{n-1} -equivariant, so defines a diffeomorphism $\hat{t}_h: T^{n-1} \times [-1, 1] \rightarrow T^{n-1} \times [-1, 1]$ which is the identity in a neighbourhood of $T^{n-1} \times \{-1, 1\}$. We now define $t_h: M \rightarrow M$ as the diffeomorphism of M such that $t_h|_U = \hat{t}_h$, $t_h|_{M \setminus U} = \text{Id}_{M \setminus U}$.

Next we show how Dehn twists can be used to give elements of infinite order in the outer automorphism group of graph manifolds.

LEMMA 5.4. *Let M be a graph manifold, with $G = \pi_1(M)$. Assume V_1, V_2 are adjacent pieces of M glued together along a common toric component H , with $G_i := \pi_1(V_i)$ and $K := \pi_1(H)$. Let $F_i \leq G_i$ be the subgroups corresponding to the fibers in V_i , and set $F = F_1 \cdot F_2 \leq K$ to be the subgroup generated by the two fiber subgroups. If $h \in K$ is chosen so that $\langle h \rangle \cap F = \{e\}$, then we have that the associated Dehn twist $\varphi := t_h$ has infinite order in $\text{Out}(G)$.*

PROOF. Suppose by way of contradiction that for some $k \geq 1$ the automorphism φ^k is equal to an internal automorphism of G , *i.e.* that there exists $\bar{g} \in G$ such that $\varphi^k(g) = \bar{g}g\bar{g}^{-1}$ for every $g \in G$. We have in particular $\bar{g}g\bar{g}^{-1} = g$ for every $g \in G_1$. By Lemma 2.14-(3), this implies that \bar{g} belongs to G_1 , whence to the center of G_1 , which coincides with the fiber subgroup F_1 of G_1 (see Remark 2.9). We conclude the conjugating element \bar{g} satisfies $\bar{g} \in F_1$.

Similarly, for every $g \in G_2$ we have $\bar{g}g\bar{g}^{-1} = h^kgh^{-k}$. Rewriting, we obtain $(h^{-k}\bar{g})g(h^{-k}\bar{g})^{-1} = g$, forcing $h^{-k}\bar{g}$ to lie in the fiber subgroup F_2 of G_2 , and hence

$$h^{-k} \in \bar{g}^{-1} \cdot F_2 \subset F_1 \cdot F_2 = F.$$

But this contradicts the fact that $\langle h \rangle \cap F = \{e\}$. We conclude that φ^k is not internal for every $k \geq 1$, as desired. \square

It is clear that the group automorphism described in Remark 5.2 is induced by a Dehn twist. As a result, Dehn twists arise naturally as basic ingredients when trying to “patch together” diffeomorphisms $\psi_i : V_i \rightarrow V'_i$ between individual pieces into a globally defined diffeomorphism $\psi : M \rightarrow M'$.

5.3. Concluding the proof of smooth rigidity

Coming back to the construction of $\psi : M \rightarrow M'$, let us consider a piece V_i of M , a component H_i of ∂V_i , and let V_j be the piece of M adjacent to V_i along H_i (we allow the case $i = j$). Denote by H_j the component of V_j which is identified to H_i in M , and by $H \subseteq M$ the image of H_i and H_j in M . We fix identifications of $\pi_1(H_i)$ with a subgroup K_i of $\pi_1(V_i)$ and of $\pi_1(H_j)$ with a subgroup K_j of $\pi_1(V_j)$ (this amounts to choosing a basepoint in H and paths joining this basepoint with the fixed basepoints of V_i and V_j). Via the fixed identifications of $\pi_1(V_i)$ and $\pi_1(V_j)$ with subgroups of $\pi_1(M)$, the groups K_i and K_j are identified with conjugated subgroups of $\pi_1(M)$, and this implies that the subgroups $\varphi(K_i)$, $\varphi(K_j)$ are conjugated in $\pi_1(M')$. By Lemma 2.14-(2), this implies that $\psi_i(H_i) = H'_i$ is glued in M' to $\psi_j(H_j) = H'_j$.

Denote by $\alpha : H_i \rightarrow H_j$ and $\alpha' : H'_i \rightarrow H'_j$ the gluing maps which enter into the definition of M and M' . We now show that the diagram

$$(5.1) \quad \begin{array}{ccc} H_i & \xrightarrow{\psi_i} & H'_i \\ \downarrow \alpha & & \downarrow \alpha' \\ H_j & \xrightarrow{\psi_j} & H'_j \end{array}$$

commutes, up to homotopy. In fact, recall that there exist $g_i, g_j \in \pi_1(M')$ such that $(\psi_i)_*(g) = g_i^{-1}\varphi(g)g_i$ for every $g \in H_i$, $(\psi_j)_*(g) = g_j^{-1}\varphi(g)g_j$ for every $g \in H_j$. Moreover, we can choose identifications $\pi_1(H'_i) \cong K'_i < g_i\pi_1(V'_i)g_i^{-1}$, $\pi_1(H'_j) \cong K'_j < g_j\pi_1(V'_j)g_j^{-1}$ in such a way that the isomorphisms $\alpha_* : K_i \rightarrow K_j$, $\alpha'_* : K'_i \rightarrow K'_j$ are induced by conjugations by an element of $\pi_1(M)$, $\pi_1(M')$ respectively.

It follows that there exists $h \in \pi_1(M')$ such that $\alpha'_*((\psi_i)_*(g)) = h(\psi_j)_*(\alpha_*(g))h^{-1}$ for every $g \in K_i$. By Lemma 2.14-(1), this implies that $h \in K'_j$, and this implies in turn that the diagram above commutes, up to homotopy. In order to properly define ψ , we now need to modify ψ_i and ψ_j in a neighbourhood of H_i and H_j ,

also taking care of the fact that ψ has eventually to induce the fixed isomorphism $\varphi: \pi_1(M) \rightarrow \pi_1(M')$.

Being homotopic *affine* diffeomorphisms of T^{n-1} , the diffeomorphisms $\alpha' \circ \psi_i$ and $\psi_j \circ \alpha$ are in fact isotopic, and this implies that ψ_i can be modified in a collar of H_i in order to make diagram (5.1) commute. This ensures that the maps ψ_i, ψ_j can be glued into a diffeomorphism $\widehat{\psi}: V_i \cup_\alpha V_j \rightarrow V'_i \cup_{\alpha'} V'_j$. As pointed out above, we are now granted that an element $h \in K'_j$ exists such that $\alpha'_*((\psi_i)_*(g)) = h(\psi_j)_*(\alpha_*(g))h^{-1}$ for every $g \in K_i$. Observe that h uniquely identifies an element of $\pi_1(H')$. It is now easily seen that if $\psi_0: V_i \cup_\alpha V_j \rightarrow V'_i \cup_{\alpha'} V'_j$ is obtained by composing $\widehat{\psi}$ with a Dehn twist along H' relative to h (or to $-h$), then ψ_0 induces on $\pi_1(V_i \cup_\alpha V_j)$ the restriction of φ .

We can apply the procedure just described along any boundary component of any piece of M , eventually obtaining the desired diffeomorphism $\psi: M \rightarrow M'$ inducing φ . \square

5.4. Mapping class group

Let M be a graph manifold. We recall that $\text{MCG}(M)$ is the mapping class group of M , *i.e.* the group of *homotopy* classes of diffeomorphisms of M onto itself. We also denote by $\text{Out}(\pi_1(M))$ the group of outer automorphisms of $\pi_1(M)$. Every diffeomorphism of M induces an isomorphism of $\pi_1(M)$, which is well-defined up to conjugacy. Since homotopic diffeomorphisms induce conjugate isomorphisms, there exists a well-defined map

$$\eta: \text{MCG}(M) \rightarrow \text{Out}(\pi_1(M)),$$

which is clearly a group homomorphism.

THEOREM 5.5. *The map $\eta: \text{MCG}(M) \rightarrow \text{Out}(\pi_1(M))$ is a group isomorphism.*

PROOF. The fact that M is aspherical (see Lemma 3.1) easily implies that η is injective, while surjectivity of η is just a restatement of Theorem 0.5. \square

REMARK 5.6. Remark 5.2 provides some evidence that the mapping class group of M should always be infinite: in fact, Dehn twists generate an abelian subgroup of $\text{MCG}(M)$, and with some effort one could probably show that such a subgroup is never finite.

REMARK 5.7. A celebrated result due to Waldhausen [**Wa1**] shows that Theorem 5.5 also holds in the case of classical closed 3-dimensional graph manifolds which either decompose into the union of at least two Seifert pieces, or do not consist of a single “small” Seifert manifold (for example, if $M = S^3$ then of course $\text{Out}(\pi_1(M)) = \{1\}$, while $\text{MCG}(M)$ has two elements). Observe however that Seifert pieces that are homeomorphic to the product $\Sigma \times S^1$, where Σ is a hyperbolic punctured surface, are never small.

In the case of classical graph manifolds with boundary, Theorem 5.5 still holds, provided that we replace the group $\text{Out}(\pi_1(M))$ with the group of the conjugacy classes of isomorphisms which preserve the peripheral structure of $\pi_1(M)$ (one says that an isomorphism of $\pi_1(M)$ preserves its peripheral structure if it sends the subgroup corresponding to a boundary component of M into the subgroup corresponding to a maybe different boundary component of M , up to conjugacy).

It is not difficult to show that Lemma 5.4 may be adapted to construct big abelian subgroups of $\text{Out}(\pi_1(M))$ also in the case of classical graph manifolds, so one expects that $\text{MCG}(M)$ should be infinite for generic 3-dimensional graph manifolds.

Algebraic properties: general graph manifolds

The aim of this chapter is the study fundamental groups of graph manifolds (and of their subgroups) with respect to some classical properties of abstract groups. We show that graph manifolds contain no non-trivial Kazhdan groups, establish the Tits alternative, show that these groups are co-Hopfian, and have universal exponential growth. The proofs of these results exploit the action of the fundamental group on the associated Bass-Serre tree. We refer the reader to [Se, Chapter I] for background on Bass-Serre theory.

Some further algebraic properties will be established later in Chapter 8, under the additional hypothesis that the graph manifold be irreducible (see Chapter 7).

6.1. Kazhdan subgroups

For M a graph manifold, the relatively simple graph of group structure of $\pi_1(M)$ can be used to extract information about the fundamental group. As a basic example, let us show how we can completely classify the subgroups of $\pi_1(M)$ which are Kazhdan (we refer the reader to [BdlHV] for a comprehensive introduction to Kazhdan groups). At the other extreme, one has amenable subgroups, which will be analyzed in the next section.

PROPOSITION 6.1. *Let M be an arbitrary graph manifold, and $H \leq \pi_1(M)$ an arbitrary subgroup. If H has Kazhdan's property (T), then H has to be the trivial group.*

PROOF. From Bass-Serre theory, we have that $\pi_1(M)$ acts on its Bass-Serre tree T , and hence we have an induced action of H on T . Kazhdan groups are known to have Serre's property (FA), i.e. any action on a tree has a globally fixed point (see [BdlHV, Section 2.3]). We conclude that H must fix a vertex in T , and hence can be viewed as a subgroup of a vertex stabilizer, which we can identify with the fundamental group of a piece V . Now the piece V splits topologically as a product, giving a splitting $\pi_1(V) \cong \pi_1(N) \times \mathbb{Z}^k$, where N is a non-compact, finite volume hyperbolic manifold. Looking at the image of H inside the first factor, we get an induced action of H on hyperbolic space. But any action of a Kazhdan group on hyperbolic space must have a global fixed point (see [BdlHV, Section 2.6]). Since $\pi_1(N)$ acts freely on hyperbolic space, we conclude that H must lie in the kernel of the natural projection $\pi_1(V) \rightarrow \pi_1(N)$, i.e. must be entirely contained in the \mathbb{Z}^k factor. Finally, the only subgroup of \mathbb{Z}^k that has Kazhdan's property (T) is the trivial group, concluding the proof. \square

By [BDS], there are finitely many conjugacy classes of homomorphisms from a Kazhdan group into a mapping class group. With respect to this issue, the behaviour of $\pi_1(M)$ is similar. In fact, as the homomorphic image of a Kazhdan group is Kazhdan, an immediate consequence of the previous Lemma is the following:

COROLLARY 6.2. *Let M be an arbitrary graph manifold. Then, there are no non-trivial homomorphisms from a Kazhdan group to $\pi_1(M)$.*

6.2. The Tits Alternative

We now show that the fundamental group of an arbitrary graph manifold satisfies the Tits Alternative.

PROPOSITION 6.3. *Let M be an arbitrary graph manifold, and $H \leq \pi_1(M)$ an arbitrary finitely generated subgroup. Then either:*

- H is solvable, or
- H contains a non-abelian free group.

PROOF. We start by considering the induced H -action on the Bass-Serre tree T associated to the splitting of $\pi_1(M)$. A ping-pong type argument shows that “most” groups that act on trees must contain a non-abelian free subgroup. More precisely (see e.g. Pays and Vallette [PaVa, Section 2]), if the group contains an element acting hyperbolically (i.e. leaving invariant a geodesic, on which it acts by a non-zero translation) and does *not* contain a non-abelian free subgroup, then it must either (i) preserve an axis, or (ii) fix an end of the tree. If we assume that our group H does not contain a non-abelian free group, we thus have three possibilities:

- (1) H consists solely of elliptic elements,
- (2) H is a subgroup of $\text{Stab}(\gamma)$, where $\gamma \subset T$ is a geodesic, or
- (3) H is a subgroup of $\text{Stab}(\mathcal{E})$, where \mathcal{E} is an end of T .

In each of these three cases, we need to show that H is solvable.

Let us consider the case (1). By induction on the number of generators, we first show that a finitely generated group consisting solely of elliptic elements must fix a vertex in the tree. If the group is generated by a single element, this follows immediately from the hypothesis. If the group is generated by k elements, consider the subgroup K generated by the first $k-1$ elements. By induction, K fixes a vertex v , and by hypothesis, the remaining generator g fixes some vertex w . If $g(v) = v$, then $\langle K, g \rangle$ fixes v and we are done. Similarly, if every element $k \in K$ satisfies $k(w) = w$, then $\langle K, g \rangle$ fixes w and we are done. So we may assume that $g(v) \neq v$, and that there exists an element $k \in K$ with $k(w) \neq w$. But this implies that the product gk is a hyperbolic element in the group, contradicting our assumption that all elements are elliptic. Finally, applying this result to the subgroup H , we see that H fixes a vertex in T , so can be viewed as a subgroup of the fundamental group of a piece. But the fundamental group of each piece is a linear group, so by work of Tits (the original Tits alternative), the fact that H does *not* contain a non-abelian free group forces H to be solvable.

Now consider case (2). The stabilizer of any geodesic naturally fits into a short exact sequence:

$$0 \rightarrow \text{Fix}(\gamma) \rightarrow \text{Stab}(\gamma) \rightarrow \text{Isom}(\gamma) \rightarrow 0$$

where $\text{Fix}(\gamma)$ is the subgroup which pointwise fixes γ , while $\text{Isom}(\gamma)$ is the induced action on γ (which we can identify with \mathbb{R}). The group $\text{Fix}(\gamma)$ fixes any given

edge on γ , hence can be identified with a subgroup of an edge stabilizer. But edge stabilizers are free abelian, so $\text{Fix}(\gamma)$ is also free abelian. On the other hand, the group $\text{Isom}(\gamma)$ is a subgroup of the group of simplicial automorphisms of \mathbb{R} (with the standard simplicial structure), hence is either $1, \mathbb{Z}_2, \mathbb{Z}$, or the infinite dihedral group \mathbb{D}_∞ . In all cases, we see that $\text{Isom}(\gamma)$ is (virtually) abelian. From the short exact sequence, we deduce that $\text{Stab}(\gamma)$ is solvable. Since H is a subgroup of the solvable group $\text{Stab}(\gamma)$, it is itself solvable.

Lastly, we consider case (3). To analyze this case, we consider the *relative translation length* map. Given an end \mathcal{E} of a tree T , and any pair of vertices $v, w \in T$, there are unique unit speed geodesic rays $\gamma_v, \gamma_w \subset T$ originating at v, w , and exiting into the end \mathcal{E} . One then defines the distance of the points *relative to \mathcal{E}* to be the integer $d_{\mathcal{E}}(v, w) := \lim_{t \rightarrow \infty} d(\gamma_v(t), \gamma_w(t))$. The relative translation length of an element $g \in \text{Stab}(\mathcal{E})$ is defined to be the integer $\tau(g) := \inf_v d_{\mathcal{E}}(v, g(v))$. A basic property of the relative translation length is that it defines a homomorphism $\tau : \text{Stab}(\mathcal{E}) \rightarrow \mathbb{Z}$ (see e.g. [PaVa, Lemme 4]). So our finitely generated group H fits into a short exact sequence

$$0 \rightarrow H_0 \rightarrow H \rightarrow \mathbb{Z} \rightarrow 0$$

where $H_0 = H \cap \ker(\tau)$. Since the quotient group H/H_0 is abelian, we also have the inclusion $H^{(1)} = [H, H] \leq H_0$, and hence $H^{(2)} = [H^{(1)}, H^{(1)}] \leq [H_0, H_0]$.

So to conclude the argument for case (3), it is enough to show that the group H_0 must be abelian, as the inclusion $H^{(2)} \leq [H_0, H_0] = 0$ will then yield that H is solvable. To analyze the (possibly not finitely generated) group H_0 , observe that every element in H_0 has to be elliptic. Given any two elements g, h in H_0 , they generate a finitely generated subgroup $\langle g, h \rangle$ of H_0 , which does not contain a free subgroup, and consists solely of elliptic elements. From case (1), we conclude that $\langle g, h \rangle$ must be abelian, and hence that g, h commute. Since this holds for arbitrary pairs of elements in H_0 , the latter is an abelian group, completing case (3) and the proof of the Proposition. \square

Since a group which contains a non-abelian free subgroup is automatically non-amenable, we have the immediate:

COROLLARY 6.4. *Let M be a graph manifold. If $H \leq \pi_1(M)$ is a finitely generated amenable subgroup, then H is solvable.*

6.3. Co-Hopf property

For our next result, we will need the notion of *transverse fibers*. Let M be a graph manifold, and V_1, V_2 a pair of adjacent pieces. We say that the two pieces have *transverse fibers* along a common torus T provided that, under the gluing diffeomorphism $\psi : T^+ \rightarrow T^-$ of the boundary tori, the image of the fiber subgroup of $\pi_1(T^+)$ under ψ_* intersects the fiber subgroup of $\pi_1(T^-)$ only in $\{0\}$.

THEOREM 6.5. *Let M be a graph manifold, with $\partial M = \emptyset$, and assume that M contains a pair of adjacent pieces with transverse fibers. Then the fundamental group $\pi_1(M)$ is co-Hopfian, i.e. every injective homomorphism $\phi : \pi_1(M) \hookrightarrow \pi_1(M)$ is automatically an isomorphism.*

PROOF. Let $\Gamma := \pi_1(M)$. Using ϕ , we can identify $\phi(\Gamma)$ with a subgroup of Γ , and our goal is to show the index $[\Gamma : \phi(\Gamma)]$ must be equal to one. A standard

argument shows that $[\Gamma : \phi(\Gamma)]$ must be finite, for if it wasn't, then we would have two manifold models for a $K(\Gamma, 1)$: the compact manifold M , and its non-compact cover \widehat{M} corresponding to the infinite index subgroup $\phi(\Gamma) \leq \Gamma$. Using these models to compute the top dimensional group cohomology of Γ with $\mathbb{Z}/2$ -coefficients gives:

$$\mathbb{Z}/2 = H^n(M^n; \mathbb{Z}/2) \cong H^n(\Gamma) \cong H^n(\widehat{M}; \mathbb{Z}/2) = 0,$$

a contradiction.

Now assume the index is some finite number $[\Gamma : \phi(\Gamma)] = k$, which we would like to show is equal to 1. From Theorem 4.5, we know that the image through ϕ of the subgroup of a piece is contained in (a conjugate of) a subgroup of a piece. This allows us to define a directed graph structure on the set of pieces of M , where we impose a directed edge from a piece V_i to a piece V_j if, under the map ϕ , the subgroup of the piece V_i maps into (a conjugate of) the subgroup of the piece V_j (recall from Lemma 2.14 that distinct pieces have non-conjugate fundamental groups). Note that in the resulting directed graph, each vertex has both indegree and outdegree equal to one, so the graph is just a union of finitely many directed cycles. Letting s denote the least common multiple of the lengths of all cycles, we can now consider a new map $\psi := \phi^s$ to be the s -fold composition of ϕ . We clearly have $[\Gamma : \psi(\Gamma)] = k^s$, and by the choice of s , for each piece V in M , we have the property that $\psi(\pi_1(V)) \subset \pi_1(V)$ (up to conjugation).

Let us consider the degree k^s cover $\widehat{M} \rightarrow M$ corresponding to the subgroup $\psi(\Gamma)$. The manifold \widehat{M} naturally inherits the structure of a graph manifold, by simply looking at all lifts of the pieces of M . Since ψ gives us an isomorphism of fundamental groups $\pi_1(M) = \Gamma \cong \psi(\Gamma) = \pi_1(\widehat{M})$, this implies that \widehat{M} has the same number of pieces as M (see Theorem 4.4). This can only happen if, under our covering map, each piece of M lifts to a single piece in \widehat{M} . We conclude that for each piece V , we must have $[\pi_1(V) : \psi(\pi_1(V))] = k^s$. Let us now fix an arbitrary piece V in M , and let V be homeomorphic to $N \times T^d$, where as usual N is a non-compact finite volume hyperbolic manifold of dimension ≥ 3 and T^d is a d -dimensional torus. The group $\Lambda := \pi_1(V)$ is isomorphic to $\pi_1(N) \times \mathbb{Z}^d$, and ψ restricts to give us an injective map from this group to itself.

We now analyze the possible injective maps from $\Lambda = \pi_1(N) \times \mathbb{Z}^d$ into itself (this is similar to the analysis in Lemma 5.1). Let $\rho : \Lambda \rightarrow \pi_1(N)$ be the natural projection onto the $\pi_1(N)$ factor. As a first step, we consider the effect of ψ on the \mathbb{Z}^d factor in Λ , and show that its image must be contained in the \mathbb{Z}^d factor. Look at the image of $\psi(\mathbb{Z}^d) \leq \Lambda$ under the ρ map. The group $\rho(\psi(\mathbb{Z}^d))$ is a free abelian subgroup of $\pi_1(N)$, and our goal is to show it is trivial. Since the \mathbb{Z}^d factor is the center of the group Λ (see Remark 2.9), we see that all of $\rho(\psi(\Lambda))$ is contained in the centralizer of $\rho(\psi(\mathbb{Z}^d))$. But inside the group $\pi_1(N)$, the centralizer of any *non-trivial* free abelian subgroup is itself free abelian (see Lemma 2.8). This implies that $\rho(\psi(\mathbb{Z}^d))$ is indeed trivial, because otherwise the preimage of its centralizer under ρ should also be free abelian, but should contain an embedded copy $\psi(\Lambda)$ of the non-abelian group Λ . Since $\rho(\psi(\mathbb{Z}^d))$ is indeed trivial, we conclude that $\psi(\mathbb{Z}^d) \leq \ker(\rho) \cong \mathbb{Z}^d$. In other words, we have just established that the map ψ embeds the \mathbb{Z}^d factor into itself.

Next, let us see how the map ψ behaves on the $\pi_1(N)$ factor, by again considering the composition with ρ . From the discussion in the previous paragraph, we have that $\rho(\psi(\Lambda)) = \rho(\psi(\pi_1(N)))$. Since $\psi(\Lambda)$ has finite index in Λ , the same holds for

any homomorphic image, giving us that $\rho(\psi(\pi_1(N)))$ has finite index in $\pi_1(N)$. But the group $\pi_1(N)$ is known to be *cofinitely Hopfian* (see [BGHM, Prop. 4.2]), i.e. any homomorphism $\pi_1(N) \rightarrow \pi_1(N)$ whose image has finite index is automatically an isomorphism. We conclude that the composite $\rho \circ \psi$ maps $\pi_1(N)$ isomorphically onto $\pi_1(N)$. Summarizing our discussion so far, in terms of the two factors in the group Λ , we can decompose the morphism ψ as $\psi(g, v) = (\phi(g), \nu(g) + Lv)$, where $\phi \in \text{Aut}(\pi_1(N))$, $\nu \in \text{Hom}(\pi_1(N), \mathbb{Z}^d)$, and $L \in GL_d(\mathbb{Z})$.

To calculate the index of $\psi(\Lambda)$ in Λ , consider the automorphism $\widehat{\psi} \in \text{Aut}(\Lambda)$ defined via $\widehat{\psi}(g, v) = (g, -\nu(\phi^{-1}(g)) + v)$. An easy computation shows that $(\widehat{\psi} \circ \psi)(g, v) = (\phi(g), Lv)$, allowing us to see that the index is

$$k^s = [\Lambda : \psi(\Lambda)] = [\Lambda : \widehat{\psi}(\psi(\Lambda))] = [\mathbb{Z}^d : L(\mathbb{Z}^d)] = |\det(L)|.$$

This formalizes the statement that the degree k^s cover $\widehat{M} \rightarrow M$ comes from unfolding the torus factors in each piece of M (along with sliding the base over the fiber, which has no affect on the degree).

Finally, let us return to our manifold M , and exploit the hypothesis on transverse fibers. Let V_1, V_2 be the pair of adjacent pieces with transverse fibers along the common torus T . The torus T corresponds to a \mathbb{Z}^{n-1} subgroup of Γ , and the two pieces give splittings of this group into direct sums $F_1^s \oplus B_1^{n-1-s} = \mathbb{Z}^{n-1} = F_2^t \oplus B_2^{n-1-t}$, where F_i are the fiber subgroups and B_i are the base subgroups. The homomorphism ψ takes \mathbb{Z}^{n-1} into itself, and by the analysis above, we can compute the index in two possible ways:

$$|\det(L_1)| = [F_1 : \psi(F_1)] = [\mathbb{Z}^{n-1} : \psi(\mathbb{Z}^{n-1})] = [F_2 : \psi(F_2)] = |\det(L_2)|$$

where L_i is a matrix representing the ψ action on F_i . Therefore, we get $k^s = |\det(L_1)| = |\det(L_2)| = |\det(\widehat{L})|$, where \widehat{L} is a matrix representing the ψ action on \mathbb{Z}^{n-1} . We will now show that this forces $k^s = 1$, whence the conclusion.

Since we have transverse fibers, we have $F_1 \cap F_2 = \{0\}$. Let us denote by K the subgroup $F_1 \oplus F_2 \subseteq \mathbb{Z}^{n-1}$, and let us set $J = \{v \in \mathbb{Z}^{n-1} \mid mv \in K \text{ for some } m \in \mathbb{Z}\}$. Of course, K is a finite index subgroup of J , and the ψ -invariance of K implies that also J is ψ -invariant. Our choices also ensure that the quotient group \mathbb{Z}^{n-1}/J is free abelian. Since ψ is injective, the following equalities hold:

$$[J : K] [K : \psi(K)] = [J : \psi(K)] = [J : \psi(J)] [\psi(J) : \psi(K)] = [J : \psi(J)] [J : K].$$

This tells us that

$$[J : \psi(J)] = [K : \psi(K)] = |\det(L_1) \cdot \det(L_2)| = k^{2s}.$$

Moreover, ψ induces a homomorphism $\overline{\psi}: \mathbb{Z}^{n-1}/J \rightarrow \mathbb{Z}^{n-1}/J$, and we have of course $\det(\widehat{L}) = \det(L_J) \cdot \det(\overline{L})$, where L_J and \overline{L} are matrices representing $\psi|_J$ and $\overline{\psi}$ respectively. Since $\det(\overline{L}) \geq 1$, we finally get

$$k^{2s} = [J : \psi(J)] = |\det(L_J)| \leq |\det(\widehat{L})| = [\mathbb{Z}^{n-1} : \psi(\mathbb{Z}^{n-1})] = k^s.$$

We conclude from this inequality that $k = 1$, giving us that $[\Gamma : \phi(\Gamma)] = k = 1$, as desired. \square

REMARK 6.6. In Theorem 6.5, the assumption that M contains a pair of adjacent pieces with transverse fibers is obviously necessary. In fact, if N is any hyperbolic manifold with toric cusps and $d \geq 1$, then the fundamental groups of

the graph manifolds $\overline{N} \times T^d$ and $D\overline{N} \times T^d$, where $D\overline{N}$ is the double of \overline{N} , are not co-Hopfian.

REMARK 6.7. Most arguments proving that the fundamental group of a closed manifold is co-Hopfian usually involve invariants which are multiplicative under coverings. Two such invariants which are commonly used are the Euler characteristic χ , and the simplicial volume. But in the case where every piece in our graph manifold has non-trivial fiber (*i.e.* of dimension greater than 0), both these invariants vanish. In fact, if $V = \overline{N} \times T^d$, $d > 0$, is a piece with non-trivial fiber, then $\chi(V) = \chi(\overline{N}) \times \chi(T^d) = 0$. Moreover, the pair $(V, \partial V)$ admits a self-map of degree greater than one, and this easily implies that the (relative) simplicial volume of V vanishes. Suppose now that a compact manifold M is obtained by gluing a (maybe disconnected) M' along a pair of π_1 -injective toric boundary components. Since the Euler characteristic of the torus is zero we have $\chi(M) = \chi(M')$, while the amenability of \mathbb{Z}^d and the additivity results proved in [Ku] imply that the (relative) simplicial volumes of M and M' coincide. Together with an obvious inductive argument, this implies that $\chi(M) = \|M\| = 0$ for a graph manifold, provided all its pieces have non-trivial fibers.

Conversely, the (relative) simplicial volume is additive with respect to gluings along π_1 -injective tori, and it never vanishes on a cusped hyperbolic manifold. So if there is a single piece in M which is purely hyperbolic (*i.e.* has trivial fiber), then $\|M\| > 0$. Similarly, the Euler characteristic of an even dimensional cusped hyperbolic manifold is never zero, so a similar conclusion holds. We summarize this discussion in the following:

PROPOSITION 6.8. *Let M be a graph manifold. Then*

- (1) $\|M\| = 0$ if and only if every piece in M has non-trivial fibers, and
- (2) if M is even dimensional, then $\chi(M) = 0$ if and only if every piece of M has non-trivial fibers.

6.4. Uniformly exponential growth

We now consider the notion of *growth* of a group G . Fixing a finite, symmetric generating set Σ , one considers the Cayley graph $C_\Sigma(G)$ of G with respect to the generating set Σ . The graph $C_\Sigma(G)$ is viewed as a metric space by setting every edge to have length = 1. For any positive real number r , we can look at the ball of radius r in $C_\Sigma(G)$ centered at the identity element, and let $N_\Sigma(r)$ count the number of vertices lying within that ball. The group has *exponential growth* provided there exists a real number $\lambda_\Sigma > 1$ with the property $N_\Sigma(r) \geq \lambda_\Sigma^r$. The property of having exponential growth is a quasi-isometry invariant, hence does not depend on the choice of generating set Σ , though the specific constant λ_Σ does depend on the choice of generating set. It is easy to see that any group which contains a free subgroup (such as the fundamental groups of our graph manifolds) automatically has exponential growth. The more sophisticated notion of *uniform exponential growth* has been the subject of recent work. A group G has uniform exponential growth if there exists a $\lambda > 1$ with the property that, for every finite symmetric generating set Σ , we have $N_\Sigma(r) \geq \lambda^r$. The point here is that the constant λ is independent of the generating set Σ . Non-elementary Gromov hyperbolic groups are known to have uniform exponential growth (see Koubi [Ko]), while CAT(0)

groups might not even have exponential growth (as the example of \mathbb{Z}^n shows). In our situation, an easy argument shows:

PROPOSITION 6.9. *If M is an arbitrary graph manifold, then $\pi_1(M)$ has uniform exponential growth.*

PROOF. Bucher and de la Harpe [**Bu-dlH**] have analyzed uniform exponential growth for groups which split as an amalgam (or as an HNN extension). It follows immediately from their work that if the graph of group description of $\pi_1(M)$ does *not* reduce to a single vertex, then $\pi_1(M)$ has uniform exponential growth. So we merely need to consider the remaining case, where M has a single piece. In this case, $\pi_1(M)$ splits as a product $\pi_1(V) \times \mathbb{Z}^k$, where V is a non-compact, finite volume hyperbolic manifold. But projecting onto the first factor, we see that $\pi_1(M)$ surjects onto a group of uniform exponential growth (by work of Eskin, Mozes, and Oh [**EMO**]). It follows that $\pi_1(M)$ also has uniform exponential growth, concluding the proof of the Lemma. \square

Recall that given a Riemannian metric g on a compact manifold M , the *volume growth entropy* of the metric is defined to be the limit

$$h_{vol}(M, g) := \lim_{r \rightarrow \infty} \frac{1}{r} \log (Vol_{\tilde{g}}(B(r)))$$

where $B(r)$ is the ball of radius r centered at a fixed point in the universal cover (\tilde{M}, \tilde{g}) with the pull-back metric from (M, g) . Work of Manning [**Ma**] shows that the *topological entropy* $h_{top}(M, g)$ of the geodesic flow on the unit tangent bundle of M satisfies the inequality $h_{top}(M, g) \geq h_{vol}(M, g)$. An immediate consequence of uniform exponential growth is the:

COROLLARY 6.10. *For M a graph manifold, there exists a real number $\delta_M > 0$ with the property that for any Riemannian metric g on M , normalized to have diameter = 1, we have the inequality $h_{top}(M, g) \geq h_{vol}(M, g) \geq \delta_M > 0$.*

Part 2

Irreducible graph manifolds: coarse geometric properties

Irreducible graph manifolds

In order to study in detail the quasi-isometric properties of the fundamental groups of graph manifolds, we would like to find conditions which ensure that walls and chambers are quasi-isometrically embedded in the universal covering. In order to do this, we introduce below the notion of *irreducible* graph manifolds. As usual, we denote by M a n -dimensional graph manifold with universal covering \widetilde{M} .

Recall that if $C \subseteq \widetilde{M}$ is a chamber, then there exists an isometry $\varphi: C \rightarrow B \times \mathbb{R}^k$, where $B \subseteq \mathbb{H}^{n-k}$ is a neutered space (such an isometry is unique up to postcomposition with the product of isometries of B and \mathbb{R}^k). Also recall that B is the *base* of C , and $F = \mathbb{R}^k$ the *fiber* of C . If $\pi_B: C \rightarrow B$, $\pi_F: C \rightarrow \mathbb{R}^k$ are the natural projections, for every $x, y \in C$, we denote by $d_B(x, y)$ the distance (with respect to the path metric of B) between $\pi_B(x)$ and $\pi_B(y)$, and by $d_F(x, y)$ the distance between $\pi_F(x)$ and $\pi_F(y)$ (so by construction $d_C^2 = d_B^2 + d_F^2$).

Let M be a graph manifold which decomposes into the pieces V_1, \dots, V_k , and let T be a boundary torus of some V_i of M . Recall that, if $p_i: V_i \rightarrow \overline{N}_i$ is the natural projection, then the fiber subgroup of T is the kernel of the map $\pi_1(T) \rightarrow \pi_1(\overline{N}_i)$. Let us denote by T_l^+, T_l^- , $l = 1, \dots, r$ the paired boundary tori of the pieces V_i 's, in such a way that M is obtained by gluing the V_i 's along the affine diffeomorphisms $\psi_l: T_l^+ \rightarrow T_l^-$, $l = 1, \dots, r$. Recall that the gluing diffeomorphism $\psi_l: T_l^+ \rightarrow T_l^-$ is *transverse* if the image of the fiber subgroup of $\pi_1(T_l^+)$ via $(\psi_l)_*$ intersects the fiber subgroup of $\pi_1(T_l^-)$ in the identity.

DEFINITION 7.1. The graph manifold M is *irreducible* if every gluing diffeomorphism $\psi_l: T_l^+ \rightarrow T_l^-$ is transverse.

Since all the groups involved in the definition of transversality are abelian, we can forget about basepoints, so transverse gluings and irreducibility are indeed well-defined. Simple examples of irreducible graph manifolds include the doubles of truncated finite volume hyperbolic manifolds with toric cusps, as well as twisted doubles of such manifolds (in the sense of [ArFa]).

The boundary of each internal wall W of \widetilde{M} decomposes into the union of two connected components W_+, W_- , while if W is a boundary wall, we simply set $W_+ = W_- = W$. We call W_+, W_- the *thin walls* associated to W , and we denote by d_{W_\pm} the path metric on W_\pm induced by the restriction of the Riemannian structure of \widetilde{M} . If W is an internal wall, then the canonical product structure on the image of W in M induces a canonical product structure $W = \mathbb{R}^{n-1} \times [-3, 3]$ with $W_\pm = \mathbb{R}^{n-1} \times \{\pm 3\}$. If $p = (x, 3) \in W_+$, $q = (y, -3) \in W_-$, we say that p, q are *tied* to each other if and only if $x = y$. If W is a boundary wall, we say that $p \in W_+ = W$ is tied to $q \in W_- = W$ if and only if $p = q$. Finally, for every wall W we denote by $s_W: W_+ \rightarrow W_-$ the map that associates to each $p \in W_+$ the point

$s_W(p) \in W_-$ tied to p . Note that, by the restriction on our gluing maps, the map s_W is an affine diffeomorphism.

LEMMA 7.2. *For W an arbitrary wall, we have:*

- (1) *if C is the chamber containing W_\pm , then the inclusion $(W_\pm, d_{W_\pm}) \hookrightarrow (C, d_C)$ is isometric;*
 - (2) *the inclusion $(W_\pm, d_{W_\pm}) \rightarrow (W, d_W)$ is a biLipschitz embedding and a quasi-isometry;*
 - (3) *the map $s_W: (W_+, d_{W_+}) \rightarrow (W_-, d_{W_-})$ is a biLipschitz homeomorphism.*
- (Of course, points (2) and (3) are trivial if W is a boundary wall).*

PROOF. By construction, if W_\pm is a thin wall and C is the chamber containing W_\pm , then W_\pm is a totally geodesic (in the metric sense) hypersurface of C . In particular, the path metric induced on W_\pm by the Riemannian structure on \widetilde{M} is isometric to the restriction of d_C , whence (1).

Concerning (2), first observe that, by definition of induced path metric, the inclusion $i: W_\pm \hookrightarrow W$ is 1-Lipschitz. The map i is the lift of an embedding which induces an isomorphism on fundamental groups, so by the Milnor-Svarc Lemma, i is a quasi-isometry. This guarantees that i is biLipschitz at large scales, *i.e.* that there exist constants $C' \geq 1$, $R > 0$ such that

$$d_{W_\pm}(x, y) \leq C' d_W(x, y) \quad \text{whenever } d_{W_\pm}(x, y) \geq R.$$

We need to control distances within the range $0 \leq d_{W_\pm}(x, y) \leq R$. Observe that this inequality describes a region $K \subseteq W_\pm \times W_\pm$ which is invariant under the obvious diagonal \mathbb{Z}^{n-1} -action. Moreover, the quotient space K/\mathbb{Z}^{n-1} is easily seen to be compact. If $K' = K \setminus \{(x, x), x \in W_\pm\}$, then the ratio d_{W_\pm}/d_W defines a positive continuous function on K' . It is not difficult to see that such a function extends to a continuous $f: K \rightarrow \mathbb{R}$ such that $f(x, x) = 1$ for every $x \in W_\pm$. Moreover, f is obviously \mathbb{Z}^{n-1} -equivariant, so compactness of K/\mathbb{Z}^{n-1} implies that f is bounded above by some constant C'' . This implies that i is $\max\{C', C''\}$ -biLipschitz, giving (2).

Similarly, s_W is obtained by lifting to the universal coverings a diffeomorphism between compact manifolds, and is therefore biLipschitz. \square

The following lemma shows how irreducibility is related to the behaviour of the metric of \widetilde{M} near the internal walls. Informally, it shows that points which almost lie on the same fiber of a thin wall are tied to points that are forced to lie on distant fibers of the adjacent chamber.

LEMMA 7.3. *Suppose $\psi_l: T_l^+ \rightarrow T_l^-$ is transverse. Let $W \subseteq \widetilde{M}$ be a (necessarily internal) wall projecting to a regular neighbourhood of $T_l^+ = T_l^-$ in M , and let $C_+, C_- \subseteq \widetilde{M}$ be the chambers adjacent to W with bases B_+, B_- . Then there exists $k \geq 1$ such that the following holds: let $x_+, y_+ \in W \cap C_+ = W_+$ (resp. $x_-, y_- \in W \cap C_- = W_-$) be such that x_+ is tied to x_- and y_+ is tied to y_- ; then*

$$d_{C_+}(x_+, y_+) \geq k d_{B_+}(x_+, y_+) \implies d_{C_-}(x_-, y_-) \leq k d_{B_-}(x_-, y_-).$$

PROOF. Suppose by contradiction that there exist sequences $\{x_+^n\}, \{y_+^n\}$ of points in W_+ such that

$$(7.1) \quad d_{C_+}(x_+^n, y_+^n) > n d_{B_+}(x_+^n, y_+^n), \quad d_{C_-}(x_-^n, y_-^n) > n d_{B_-}(x_-^n, y_-^n).$$

Recall that W_+ and W_- are endowed with a canonical affine structure, and the map $s_W: W_+ \rightarrow W_-$ defined before Lemma 7.2 is an affine diffeomorphism. Let Z_+, Z_- be the vector spaces underlying the affine spaces W_+, W_- , and denote by $\widehat{s}_W: Z_+ \rightarrow Z_-$ the linear map associated to s_W .

The product decompositions of $C_+ = B_+ \times F_+$ and $C_- = B_- \times F_-$ induce direct sum decompositions

$$Z_+ = \widehat{B}_+ \oplus \widehat{F}_+, \quad Z_- = \widehat{B}_- \oplus \widehat{F}_-,$$

and transversality of ψ_l implies that $\widehat{s}_W(\widehat{F}_+) \cap \widehat{F}_- = \{0\}$.

For every $n \in \mathbb{N}$, we denote by $v_+^n \in \widehat{F}_+, w_+^n \in \widehat{B}_+$ (resp. $v_-^n \in \widehat{F}_-, w_-^n \in \widehat{B}_-$) the vectors uniquely determined by the conditions $y_+^n - x_+^n = v_+^n + w_+^n, y_-^n - x_-^n = v_-^n + w_-^n$.

By Lemma 7.2-(1), the restrictions of the distances d_{C_+} and d_{C_-} to W_+ and W_- are induced by Euclidean norms $\|\cdot\|_+, \|\cdot\|_-$ on Z_+, Z_- . The inequalities (7.1) may now be rewritten in the following way:

$$(7.2) \quad \frac{\|v_+^n + w_+^n\|_+}{n} > \|w_+^n\|_-, \quad \frac{\|v_-^n + w_-^n\|_-}{n} > \|w_-^n\|_-.$$

Up to rescaling, we may suppose that $\|v_+^n + w_+^n\|_+ = 1$ for every n . Since s_W is biLipschitz, there exists $\alpha \geq 1$ such that $\alpha^{-1} \leq \|v_-^n + w_-^n\|_- \leq \alpha$ for every n . In particular, up to passing to subsequences, we may suppose that the sequences $\{v_+^n\}, \{w_+^n\}, \{v_-^n\}, \{w_-^n\}$ converge to $v_+ \in \widehat{F}_+, w_+ \in \widehat{B}_+, v_- \in \widehat{F}_-, w_- \in \widehat{B}_-$. Moreover, we have $\widehat{s}_W(v_+ + w_+) = v_- + w_-$. As n tends to infinity, inequalities (7.2) imply $w_+ = 0, w_- = 0$, so $\widehat{s}_W(v_+) = v_-$. Since $\|v_+\|_+ = \|v_+ + w_+\|_+ = 1$, we have that $\widehat{s}_W(v_+) = v_-$ is a non-trivial element in $\widehat{s}_W(\widehat{F}_+) \cap \widehat{F}_- = \{0\}$, and this provides the desired contradiction. \square

We now come to the proof of the main result of this section, which states that if M is irreducible then the inclusion of walls and chambers in \widetilde{M} are quasi-isometric.

The following Proposition is inspired by [Os, Lemma 3.2]:

PROPOSITION 7.4. *Let B be a neutered space. Then there exists a constant Q only depending on B such that the following result holds. Let $\gamma \subseteq B$ be a loop obtained by concatenating a finite number of paths $\alpha_1, \gamma_1, \dots, \alpha_n, \gamma_n$, where*

- each α_i is a geodesic on a horosphere $O_i \subseteq \partial B$, and
- each γ_i is any path in B connecting the endpoint of α_i with the starting point of α_{i+1} .

Let $D \subseteq \{1, \dots, n\}$ be a distinguished subset of indices such that $O_h \neq O_i$ for every $h \in D, i \in \{1, \dots, n\}, i \neq h$. Then

$$\sum_{h \in D} L(\alpha_h) \leq Q \sum_{i=1}^n L(\gamma_i).$$

PROOF. Let B be a neutered space, and recall that by the very definitions, the group of isometries of B contains a discrete torsion-free cocompact subgroup Γ . The quotient $N = B/\Gamma$ is obtained by removing horospherical neighbourhoods of the cusps from a finite-volume hyperbolic manifold. As a consequence, there exists $R > 0$ such that the distance between every pair of distinct connected components




FIGURE 1. Proposition 7.4 provides a bound on the lengths of the α_i 's in terms of the lengths of the γ_i 's.

of ∂B is at least R , so that

$$(7.3) \quad n \leq \frac{\sum_j L(\gamma_j)}{R}.$$

Let $\{H_1, \dots, H_l\}$ be the collection of subgroups of Γ obtained by choosing a representative in each conjugacy class of cusp subgroups of N . Recall that Γ is relatively hyperbolic with respect to the H_i 's (see Section 9.6 for a definition of relative hyperbolicity, and *e.g.* [Fa2] for a proof that Γ is relatively hyperbolic with respect to the H_i 's). Choose X to be a symmetric set of generators for Γ satisfying the assumptions of [Os, Lemma 3.2], and let us denote by \mathcal{C}_Γ the corresponding Cayley graph of Γ with distance d_Γ .

We denote by $\bar{\mathcal{C}}_\Gamma$ the Cayley graph of Γ with respect to the (infinite) set of generators $(X \cup (H_1 \cup \dots \cup H_l)) \setminus \{1\}$, and by \bar{d}_Γ the path distance on $\bar{\mathcal{C}}_\Gamma$ (see [Os]).

More precisely, if \tilde{X} is a copy of X , \tilde{H}_λ is a copy of H_λ and $\mathcal{H} = \bigsqcup_{\lambda=1}^l (\tilde{H}_\lambda \setminus \{1\})$, then $\bar{\mathcal{C}}_\Gamma$ is the graph having Γ as set of vertices and $\Gamma \times (\tilde{X} \cup \mathcal{H})$ as set of edges, where if $\bar{y} \in \Gamma$ is the element corresponding to $y \in \tilde{X} \cup \mathcal{H}$, then the edge (g, y) has g and $g \cdot \bar{y}$ as endpoints. We label the edge (g, y) by the symbol y . Note that different labels may represent the same right multiplication in Γ : for instance, this is the case if there exist letters $x \in \tilde{X}$ and $y \in \tilde{H}$ representing the same element $\bar{x} = \bar{y}$ in Γ , *i.e.* if $X \cap \left(\bigcup_{\lambda=1}^l H_\lambda\right) \neq \emptyset$.

Notice that by the very definitions we have a natural inclusion $\mathcal{C}_\Gamma \hookrightarrow \bar{\mathcal{C}}_\Gamma$. Let q be a (non-based) loop in $\bar{\mathcal{C}}_\Gamma$ labelled by the (cyclic) word w with letters in $\tilde{X} \cup \mathcal{H}$. Recall from [Os] that a subpath of a loop q in $\bar{\mathcal{C}}_\Gamma$ is a H_λ -subpath if it is labelled by a subword of w with letters in \tilde{H}_λ . A H_λ -component of q is a maximal H_λ -subpath of q . A H_λ -component q' of q is *not isolated* if there exists a H_λ -component $q'' \neq q'$ of q such that a vertex in q' and a vertex of q'' are joined by an edge labelled by a letter in \tilde{H}_λ (in algebraic terms this means that such vertices belong to the same left coset of H_λ in Γ).

Starting from γ , we wish to construct a loop $\bar{\gamma}$ in $\bar{\mathcal{C}}_\Gamma$. Milnor-Svarc's Lemma provides a (μ, ϵ) -quasi-isometry $\varphi: B \rightarrow \mathcal{C}_\Gamma$. Up to increasing ϵ , we can require that φ maps every point of B onto a vertex of \mathcal{C}_Γ , *i.e.* onto an element of Γ , and that

every horosphere $O \subseteq \partial B$ is taken by φ onto a lateral class of some H_λ . It is easy to see that if φ maps the horospheres $O, O' \subseteq \partial B$ onto the same lateral class of the same H_λ , then $O = O'$. Fix $i \in \{1, \dots, n\}$, suppose that γ_i is parametrized by arc length, denote by m_i the least integer number such that $L(\gamma_i) \leq m_i$, and set $p_i^j = \varphi(\gamma_i(jL(\gamma_i)/m_i)) \in \Gamma$ for $j = 0, \dots, m_i$. Due to our choices we have $p_i^0 \in \varphi(O_i)$ and $p_i^{m_i} \in \varphi(O_{i+1})$. Now let $\tilde{\gamma}_i$ be the path in \mathcal{C}_Γ obtained by concatenating the geodesics joining p_i^j and p_i^{j+1} , $j = 0, \dots, m_i - 1$, and let $\bar{\gamma}_i$ be the path in $\bar{\mathcal{C}}_\Gamma$ obtained by taking the image of $\tilde{\gamma}_i$ under the inclusion $\mathcal{C}_\Gamma \hookrightarrow \bar{\mathcal{C}}_\Gamma$. Observe that by construction every edge of $\bar{\gamma}_i$ is labelled by a symbol in \tilde{X} , so no $\bar{\gamma}_i$ contains any H_λ -subpath.

As $m_i \leq L(\gamma_i) + 1$ by our choice of m_i , we have the estimate:

$$L(\bar{\gamma}_i) = L(\tilde{\gamma}_i) = \sum_{j=0}^{m_i-1} d_\Gamma(p_i^j, p_i^{j+1}) \leq \mu L(\gamma_i) + m_i \epsilon \leq (\mu + \epsilon) L(\gamma_i) + \epsilon.$$

Next, observe that $p_{i-1}^{m_{i-1}}$ and p_i^0 both lie on $\varphi(O_i)$, and hence belong to the same left coset of some $H_{\psi(i)}$, $\psi(i) \in \{1, \dots, l\}$. Thus we can connect $p_{i-1}^{m_{i-1}}$ and p_i^0 in $\bar{\mathcal{C}}_\Gamma$ by a path $\bar{\alpha}_i$ which is either constant (if $p_{i-1}^{m_{i-1}} = p_i^0$), or consists of a single edge labelled by a symbol in $\tilde{H}_{\psi(i)}$. Now define the loop $\bar{\gamma} = \bar{\alpha}_1 * \bar{\gamma}_1 * \dots * \bar{\alpha}_n * \bar{\gamma}_n$ in $\bar{\mathcal{C}}_\Gamma$. Using (7.3), we obtain

$$\begin{aligned} L(\bar{\gamma}) &\leq \left(\sum_{i=1}^n L(\bar{\gamma}_i) \right) + n \\ &\leq (\mu + \epsilon) \sum_{i=1}^n L(\gamma_i) + n\epsilon + n \\ &\leq \left(\mu + \epsilon + \frac{\epsilon + 1}{R} \right) \sum_{i=1}^n L(\gamma_i). \end{aligned}$$

Moreover, due to our assumption on D , for every $h \in D$ the subpath $\bar{\alpha}_h$ is (either constant or) an isolated component of $\bar{\gamma}$, so by [Os, Lemma 3.2] there exists Q' only depending on (the Cayley graphs \mathcal{C}_Γ and $\bar{\mathcal{C}}_\Gamma$ of) Γ such that for every $\lambda = 1, \dots, l$

$$\sum_{h \in D \cap \psi^{-1}(\lambda)} d_\Gamma(p_{h-1}^{m_{h-1}}, p_{m_h}^0) \leq Q' L(\bar{\gamma}),$$

whence

$$(7.4) \quad \sum_{h \in D} d_\Gamma(p_{h-1}^{m_{h-1}}, p_{m_h}^0) \leq lQ' L(\bar{\gamma}) \leq lQ' \left(\mu + \epsilon + \frac{\epsilon + 1}{R} \right) \sum_{i=1}^n L(\gamma_i).$$

On the other hand we have

$$(7.5) \quad \begin{aligned} \sum_{h \in D} d_\Gamma(p_{h-1}^{m_{h-1}}, p_{m_h}^0) &\geq \frac{1}{\mu} \sum_{h \in D} L(\alpha_h) - \epsilon n \\ &\geq \frac{1}{\mu} \sum_{h \in D} L(\alpha_h) - \frac{\epsilon}{R} \sum_{i=1}^n L(\gamma_i). \end{aligned}$$

Putting together inequalities (7.4) and (7.5) we finally get that the inequality of the statement holds for some Q only depending on μ, ϵ, Q', R . \square

Let us fix the graph manifold M which we are studying. We will now introduce various constants, which will be extensively used in the rest of the arguments for this section. Fix the following quantities:

- the constant Q : chosen so that Proposition 7.4 holds for all the bases of the chambers of \widetilde{M} .
- the constant R : the minimal distance between pairs of thin walls *not* associated to the same internal wall (note that R is also the minimal distance between pairs of removed horoballs in the bases of the chambers of \widetilde{M}).
- the constant D : the maximal distance between pairs of *tied* points on adjacent thin walls (here we refer to the path distance of the corresponding wall).
- the constant k : chosen so that Lemma 7.3 holds for all the internal walls in \widetilde{M} .
- the constant k' : chosen so that $s_W: W_+ \rightarrow W_-$ is k' -biLipschitz for every internal wall W of \widetilde{M} .
- the constant c : chosen so that all the inclusions $W_\pm \hookrightarrow W$ are c -biLipschitz (see Lemma 7.2).

These constants only depend on the geometry of M . In what follows, we will also assume without loss of generality that $Q \geq 2$ and $k \geq \sqrt{2}$.

In order to prove that walls and chambers are quasi-isometrically embedded in \widetilde{M} , we need to show that the distance between points in the same chamber can be bounded from below by the distance of the projections of the points on the base of the chamber. We begin with the following:


DEFINITION 7.5. Let W_\pm be a thin wall, take $x, y \in W_\pm$ and let γ be a continuous path joining x and y . We say that γ *does not backtrack* on W_\pm if γ intersects the wall containing W_\pm only in its endpoints.

LEMMA 7.6. *Let x, y be points on the same thin wall W_\pm and let γ be a path in \widetilde{M} which joins x to y and does not backtrack on W_\pm . If C is the chamber containing W_\pm and B is the base of C , then $L(\gamma) \geq d_B(x, y)/Q$.*

PROOF. An easy transversality argument shows that it is not restrictive to assume that the intersection of γ with C consists of a finite number of subpaths of γ . Now the sum of the lengths of such subpaths is greater than the sum of the lengths of their projections on B , which is in turn greater than $d_B(x, y)/Q$ by Proposition 7.4. \square

If the distance of two points on a thin wall is not suitably bounded by the distance of their projections on the base of the chamber they belong to, then Lemma 7.6 does not give an effective estimate. Together with Lemma 7.3, the following result shows how in this case, irreducibility allows us to “pass to the adjacent chamber” in order to obtain a better estimate.

LEMMA 7.7. *Let $x_+, y_+ \in W_+$ be points on a thin wall, let C_+ be the chamber containing W_+ , and suppose that γ is a rectifiable path joining x_+ and y_+ and*

FIGURE 2. Decomposing γ in the proof of Lemma 7.7.

intersecting C_+ only in its endpoints. Let also $x_-, y_- \in W_-$ be the points tied to x_+, y_+ , and C_- be the chamber containing x_-, y_- . Then

$$L(\gamma) \geq \frac{d_{B_-}(x_-, y_-)}{cQ} - \frac{2D}{Q},$$

where B_- is the base of C_- .

PROOF. An easy transversality argument shows that it is not restrictive to assume that γ intersects the thin walls of \widetilde{M} only in a finite number of points. Then our assumptions imply that γ decomposes as a concatenation of curves

$$\gamma = \gamma'_1 * \gamma''_1 * \gamma'_2 * \dots * \gamma''_n * \gamma'_{n+1}$$

such that γ'_i is supported in W and γ''_i has endpoints $a_i, b_i \in W_-$ and does not backtrack on W_- for every i (see Figure 2). Let us suppose $n \geq 1$ (the case $n = 0$ being easier). Since $d_W(x_-, x_+) \leq D$ we have

$$d_{B_-}(x_-, a_1) \leq d_{W_-}(x_-, a_1) \leq cd_W(x_-, a_1) \leq c(D + L(\gamma'_1)),$$

and analogously we get $d_{B_-}(y_-, b_n) \leq c(D + L(\gamma'_{n+1}))$. Moreover Lemma 7.6 implies $d_{B_-}(a_i, b_i) \leq Q \cdot L(\gamma''_i)$ for every $i = 1, \dots, n$, and we also have $d_{B_-}(b_i, a_{i+1}) \leq d_{W_-}(b_i, a_{i+1}) \leq cd_W(b_i, a_{i+1}) \leq cL(\gamma'_{i+1})$ for every $i = 1, \dots, n-1$. Putting together all these inequalities we finally get

$$\begin{aligned} d_{B_-}(x_-, y_-) &\leq d_{B_-}(x_-, a_1) + \sum_{i=1}^n d_{B_-}(a_i, b_i) + \sum_{i=1}^{n-1} d_{B_-}(b_i, a_{i+1}) + d_{B_-}(b_n, y_-) \\ &\leq 2cD + c \sum_{i=1}^{n+1} L(\gamma'_i) + Q \sum_{i=1}^n L(\gamma''_i) \\ &\leq 2cD + cQL(\gamma) \end{aligned}$$

whence the conclusion. \square

In order to proceed to the main argument we finally need the following lemma, which describes how to get rid of the backtracking of a geodesic.

If γ is a path and $r = \gamma(t_0)$, $s = \gamma(t_1)$, with an abuse we will denote by $[r, s]$ the subpath $\gamma|_{[t_0, t_1]}$ of γ . We say that γ is *minimal* if for every chamber C , the set $\gamma \cap \mathring{C}$ is a finite collection of paths each of which connects distinct walls of C . Moreover, γ is *good* if it is minimal and for every thin wall X contained in a chamber C there are at most 2 endpoints of paths in $\gamma \cap \mathring{C}$ belonging to X . Notice that, since chambers are uniquely geodesic and every thin wall is totally geodesic in the chamber in which it is contained, every geodesic of \widetilde{M} is minimal.


LEMMA 7.8. *There exists a constant $\beta \geq 1$ depending only on the geometry of \widetilde{M} such that the following result holds. Let x, y be points belonging to the same wall of \widetilde{M} . Then there exists a good path γ in \widetilde{M} connecting x and y such that $L(\gamma) \leq \beta d(x, y)$.*

PROOF. We first introduce some terminology. If X is a thin wall contained in the wall W , we say that a path $\theta : [t_0, t_1] \rightarrow \widetilde{M}$ is *external* to X if $\theta(t_0) \in X$, $\theta(t_1) \in X$ and $\theta|_{(t_0, t_1)}$ is supported in $\widetilde{M} \setminus W$ (this is equivalent to asking that θ does not backtrack on X , but this new terminology will prove more appropriate here). Moreover, if γ is a minimal path and n is the number of subpaths of γ external to X , we say that the *exceeding number* of γ on X is equal to $\max\{0, n - 1\}$. The exceeding number $e(\gamma)$ of γ is the sum of the exceeding numbers of γ on all the thin walls. Finally, we denote by $j(\gamma)$ the sum over *all* the chambers C of \widetilde{M} of the number of connected components of $\gamma \cap \mathring{C}$. It is readily seen that a path γ is good if and only if it is minimal and $e(\gamma) = 0$.

Let $\Delta > 0$ be a constant, chosen in such a way that every torus in M obtained as a projection of a thin wall of \widetilde{M} has diameter (with respect to its intrinsic path metric) at most $\Delta/2$. We denote by γ_0 a geodesic in \widetilde{M} connecting x and y . As observed above, γ_0 is minimal, and if γ_i is a minimal path with $e(\gamma_i) > 0$ we will now describe how to modify it in order to get a new minimal path γ_{i+1} joining x to y . The path γ_{i+1} will be constructed so as to have $j(\gamma_{i+1}) < j(\gamma_i)$ and $L(\gamma_{i+1}) \leq L(\gamma_i) + 4\Delta + 1$. By the very definitions we have $j(\gamma_0) \leq L(\gamma_0)/R = d(x, y)/R$, so after at most $d(x, y)/R$ steps we will end up with a minimal path γ which verifies either $e(\gamma) = 0$ or $j(\gamma) \leq 1$, whence again $e(\gamma) = 0$. After setting $\beta = 1 + (4\Delta + 1)/R$, such a path satisfies all the conditions required.

So let us suppose that we have some external subpaths $[p_1, p'_1], [p_2, p'_2]$ of γ_i , with $p_1, p'_1, p_2, p'_2 \in X$ for some thin wall X contained in the chamber C . Consider deck transformations g, h which leave X (and therefore C) invariant such that $d_X(g(p_2), p'_1) \leq \Delta$, $d_X(h(p'_1), g(p'_2)) \leq \Delta$, and let $q_1, q_2 \in \gamma \cap \mathring{C}$ be chosen in such a way that q_1 (resp. q_2) slightly precedes (resp. follows) p'_1 (resp. p_2) on γ_i : more precisely, we assume that $L([q_1, p'_1]) < 1/2$, $L([p_2, q_2]) < 1/2$. We define a path γ'_{i+1} as the concatenation of the following paths (see Figure 3):

- (1) the subpath $[x, q_1]$ of γ_i ,
- (2) a path $[q_1, g(q_2)]$ in \mathring{C} obtained by slightly pushing inside \mathring{C} a geodesic in X between p'_1 and $g(p_2)$, in such a way that $L([q_1, g(q_2)]) < \Delta + 1/2 + 1/2 = \Delta + 1$,
- (3) $g([q_2, p'_2])$,
- (4) a geodesic in X between $g(p'_2)$ and $h(p'_1)$,
- (5) $h([p'_1, p_2])$,
- (6) a geodesic in X between $h(p_2)$ and p'_2 ,

FIGURE 3. Replacing γ_i with γ'_{i+1} in the proof of Lemma 7.8.(7) $[p'_2, y]$,

where geodesics in X are to be considered with respect to its path metric.

Since X is isometric to \mathbb{R}^{n-1} and the deck transformations g, h act on X as translations, it follows that the distance between $h(p_2)$ and p'_2 is at most 2Δ , and this readily yields $L(\gamma'_{i+1}) \leq L(\gamma_i) + 4\Delta + 1$. Moreover, it is easily checked that $j(\gamma'_{i+1}) = j(\gamma_i) - 1$. Now, if γ'_{i+1} is minimal we set $\gamma_{i+1} = \gamma'_{i+1}$, and we are done. On the other hand, the only possible obstruction to γ'_{i+1} being minimal is that its (open) subpath with endpoints p_1 and $g(p'_2)$ may be entirely contained in \mathring{C} . In this case, since X is totally geodesic in C we can replace the subpath $[p_1, g(p'_2)]$ with a geodesic on X , thus obtaining a minimal path γ_{i+1} with $L(\gamma_{i+1}) \leq L(\gamma'_{i+1})$ and $j(\gamma_{i+1}) = j(\gamma'_{i+1}) - 1 < j(\gamma_i)$, whence the conclusion again. \square

REMARK 7.9. It is maybe worth mentioning that Lemma 7.8 also holds if M is *not* assumed to be irreducible. However, its proof relies heavily on the fact that thin walls support a *flat* metric: by Bieberbach Theorem, this ensures that a finite index subgroup of the covering transformations of \widetilde{M} preserving a thin wall acts on it as a group of translations, and this fact plays a fundamental role in the proof of Lemma 7.8. As a consequence, extending Lemma 7.8 to more general situations (*e.g.* when studying the universal covering of spaces obtained by gluing non-positively curved locally symmetric manifolds along *non-flat* boundaries) could be a challenging task.

LEMMA 7.10. *Fix a wall $W \subseteq \widetilde{M}$ and suppose that $\alpha \geq 1$ exists such that the following holds: if $x, y \in W_{\pm}$ are points joined by a good path γ in \widetilde{M} which does not backtrack on W_{\pm} , then*

$$d_C(x, y) \leq \alpha \cdot L(\gamma),$$

where C is the chamber containing x, y . Then W is biLipschitz embedded in \widetilde{M} .

PROOF. The inclusion $(W, d_W) \hookrightarrow \widetilde{M}$ is clearly 1-Lipschitz, so we have to check that d_W is linearly bounded below by the distance d on \widetilde{M} . More precisely, we have to show that there exists $\lambda \geq 1$ such that

$$(7.6) \quad d_W(p, q) \leq \lambda d(p, q) \quad \text{for all } p, q \in W.$$

Let γ be the path provided by Lemma 7.8 such that $L(\gamma) \leq \beta d(p, q)$, and let m be the number of the chambers adjacent to W whose internal parts intersect γ (so $m = 0, 1$ or 2). It is readily seen that γ splits as a concatenation

$$\gamma_1 * \gamma'_1 * \cdots * \gamma_m * \gamma'_m * \gamma_{m+1},$$

where the γ_i 's are contained in W and each γ'_i is a good path with endpoints on W_{\pm} which does not backtrack on W_{\pm} . Due to our assumptions and to the fact that W_{\pm} are totally geodesic in the chambers in which they are contained, the γ'_i 's can be replaced by curves contained in W in such a way that the total length of the curve so obtained does not exceed $\alpha \cdot L(\gamma)$. So

$$d_W(p, q) \leq \alpha L(\gamma),$$

and hence inequality (7.6) holds with $\lambda = \alpha \cdot \beta$. \square

THEOREM 7.11. *If M is irreducible and $W \subseteq \widetilde{M}$ is a wall, then the inclusion $(W, d_W) \hookrightarrow \widetilde{M}$ is a biLipschitz embedding. In particular, it is a quasi-isometric embedding. Moreover, the biLipschitz constant of the embedding only depends on the geometry of \widetilde{M} (i.e. it does not depend on the fixed wall W).*

PROOF. Take $x, y \in W_+$, let C be the chamber containing W_+ and let γ be a good path in \widetilde{M} which joins x to y and does not backtrack on W_+ . By Lemma 7.10, in order to conclude it is sufficient to show that the inequality

$$(7.7) \quad d_C(x, y) \leq \alpha \cdot L(\gamma).$$

holds for some $\alpha \geq 1$ only depending on \widetilde{M} (via the constants D, R, Q, k, k', c). We will have to analyze several different cases, and we will take α to be the maximum among the constants we will find in each case.

Let B, F be the base and the fiber of C . We first distinguish the case when the distance between x and y is controlled (up to a suitable constant factor) by $d_B(x, y)$ from the case when $d_C(x, y)$ is controlled by $d_F(x, y)$.

So let us suppose $d_C(x, y) \leq k d_B(x, y)$. In this case by Lemma 7.6 we have $L(\gamma) \geq d_B(x, y)/Q \geq d_C(x, y)/(kQ)$, so

$$d_C(x, y) \leq kQL(\gamma),$$

and we are done.

Let us now consider the other case and assume that $d_C(x, y) > k d_B(x, y)$. Since $d_C^2 = d_B^2 + d_F^2$ and $k > \sqrt{2}$ an easy computation shows that

$$d_F(x, y) > \frac{d_C(x, y)}{\sqrt{2}}, \quad d_F(x, y) > d_B(x, y).$$

Write $\gamma \cap \mathring{C} = \gamma_1 \cup \dots \cup \gamma_m$ where each $\gamma_i = (x_i, y_i)$ is a path in the (open) chamber \mathring{C} , let W_i be the wall containing y_i and x_{i+1} , and let l_i be the length of the projection of γ_i on the fiber F . Observe that since γ is minimal we have $m \leq L(\gamma)/R$. Of

course, we have $\sum l_i + \sum d_F(y_i, x_{i+1}) \geq d_F(x, y)$, so either $\sum l_i \geq d_F(x, y)/2$ or $\sum d_F(y_i, x_{i+1}) \geq d_F(x, y)/2$. In the first case we have

$$L(\gamma) \geq \sum L(\gamma_i) \geq \sum l_i \geq \frac{d_F(x, y)}{2} > \frac{d_C(x, y)}{2\sqrt{2}},$$

and we are done. Otherwise let us define

$$I_1 = \{i \in \{1, \dots, m-1\} \mid kd_B(y_i, x_{i+1}) \leq d_C(y_i, x_{i+1})\}, \quad I_2 = \{1, \dots, m-1\} \setminus I_1.$$

Since $\sum d_C(y_i, x_{i+1}) \geq \sum d_F(y_i, x_{i+1}) \geq d_F(x, y)/2$, we have two possibilities: either $\sum_{i \in I_1} d_C(y_i, x_{i+1}) \geq d_F(x, y)/4$, or $\sum_{i \in I_2} d_C(y_i, x_{i+1}) \geq d_F(x, y)/4$.

We begin by dealing with the first case. Let W_+^i be the thin wall containing x_{i+1}, y_i , denote by $x_{i+1}^- \in W_-^i$ (resp. $y_i^- \in W_-^i$) the point tied to x_{i+1} (resp. to y_i), let C_i be the chamber containing x_{i+1}^-, y_i^- , and B_i the base of C_i .

Recall that we have the estimate:

$$d_{C_i}(y_i^-, x_{i+1}^-) = d_{W_-^i}(y_i^-, x_{i+1}^-) \geq d_{W_+^i}(y_i, x_{i+1})/k' = d_C(y_i, x_{i+1})/k'.$$

So applying Lemma 7.7 and Lemma 7.3 (which gives the inequality $d_{B_i}(y_i^-, x_{i+1}^-) \geq d_{C_i}(y_i^-, x_{i+1}^-)/k$) we obtain the estimates:

$$\begin{aligned} L(\gamma) &\geq \frac{1}{cQ} \sum_{i \in I_1} d_{B_i}(y_i^-, x_{i+1}^-) - \frac{2(m-1)D}{Q} \\ &\geq \frac{1}{kcQ} \sum_{i \in I_1} d_{C_i}(y_i^-, x_{i+1}^-) - \frac{2(m-1)D}{Q} \\ &\geq \frac{1}{kcQk'} \sum_{i \in I_1} d_C(y_i, x_{i+1}) - \frac{2D}{RQ} L(\gamma) \\ &\geq \frac{1}{4kcQk'} d_F(x, y) - \frac{2D}{RQ} L(\gamma) \\ &\geq \frac{1}{4\sqrt{2}kcQk'} d_C(x, y) - \frac{2D}{RQ} L(\gamma) \end{aligned}$$

Isolating the $d_C(x, y)$ term, this gives us

$$d_C(x, y) \leq \frac{4\sqrt{2}kck'(RQ + 2D)}{R} \cdot L(\gamma).$$

which gives us the requisite estimate in the first case.

We are now left to deal with the second case, $\sum_{i \in I_2} d_C(y_i, x_{i+1}) \geq d_F(x, y)/4$. In this case we have that:

$$\sum_{i \in I_2} d_B(y_i, x_{i+1}) \geq d_F(x, y)/(4k) \geq d_C(x, y)/(4\sqrt{2}k).$$

Let γ' be the loop in C obtained by concatenating the geodesic in W_+ joining y with x , the paths of the form $\gamma \cap \dot{C}$ and the geodesics in the W_+^i 's joining y_i with x_{i+1} , and set $\bar{\gamma} = \pi_B \circ \gamma'$. If η is the sum of the lengths of the subpaths of $\bar{\gamma}$ obtained by projecting the paths in $\gamma \cap \dot{C}$ we obviously have $L(\gamma) \geq \eta$. Moreover, the properties

of γ described in Lemma 7.8 ensure that the W_i 's are pairwise distinct, and distinct from W . As such, we can apply Proposition 7.4 to $\bar{\gamma}$ thus getting

$$L(\gamma) \geq \eta \geq \frac{1}{Q} \sum_{i \in I_2} d_B(y_i, x_{i+1}),$$

whence

$$d_C(x, y) \leq 4\sqrt{2}k \sum_{i \in I_2} d_B(y_i, x_{i+1}) \leq 4\sqrt{2}kQL(\gamma).$$

This completes the last case, establishing that inequality (7.7) holds with constant

$$\alpha = \max \left\{ kQ, 2\sqrt{2}, \frac{4\sqrt{2}kck'(RQ + 2D)}{R}, 4\sqrt{2}kQ \right\} = \frac{4\sqrt{2}kck'(RQ + 2D)}{R},$$

thus proving the proposition. \square

COROLLARY 7.12. *If M is irreducible, then the inclusion of a chamber in \widetilde{M} is a biLipschitz embedding.*

PROOF. Let p, q be points in a chamber C and let δ be a geodesic of \widetilde{M} joining p to q . Then δ splits as a concatenation

$$\delta = \delta_1 * \eta_1 * \dots * \eta_n * \delta_{n+1},$$

where δ_i is a geodesic segment (with respect to the metric d on \widetilde{M}) supported in C and the endpoints p_i, q_i of η_i belong to a thin wall W_i^+ adjacent to C . By Theorem 7.11 there exists $\alpha \geq 1$ such that $d_{W_i}(p_i, q_i) \leq \alpha d(p_i, q_i)$, and this implies in turn that $d_{W_i^+}(p_i, q_i) \leq \alpha d(p_i, q_i)$, so we may replace every η_i with a path $\eta'_i \subseteq W_i^+$ having the same endpoints as η_i and length that does not exceed $\alpha d(p_i, q_i)$. The path

$$\delta' = \delta_1 * \eta'_1 * \dots * \eta'_n * \delta_{n+1}$$

is supported in C and has length at most $\alpha d(p, q)$, so $d_C(p, q) \leq \alpha d(p, q)$, and we are done. \square

COROLLARY 7.13. *Suppose that M is irreducible. Then, the inclusion of chambers, walls and fibers (with their path metrics) in \widetilde{M} are quasi-isometric embeddings. In particular:*

- *If $C \subseteq \widetilde{M}$ is a chamber, then C is quasi-isometric (with the metric induced by \widetilde{M}) to a product $B \times \mathbb{R}^k$, where B is a neutered space.*
- *If $W \subseteq \widetilde{M}$ is a wall, then W is quasi-isometric (with the metric induced by \widetilde{M}) to \mathbb{R}^{n-1} .*
- *If $F \subseteq \widetilde{M}$ is a fiber, then F is quasi-isometric (with the metric induced by \widetilde{M}) to \mathbb{R}^h , $h \leq n - 3$.*

REMARK 7.14. It would be interesting to find less restrictive conditions under which the walls of \widetilde{M} are ensured to be quasi-isometrically embedded. However, we point out that such a task could turn out to be quite difficult. In fact, even in the context of Gromov hyperbolic groups, it is not easy to decide whether vertex groups of a graph of hyperbolic groups are quasi-isometrically embedded in the fundamental group of the graph of groups. In [Ka] it is shown that this is the case under the assumption that the graph of groups is *acylindrical* (see Subsection 8.1 for

the definition) and edge groups are quasi-isometrically embedded in vertex groups (this last condition is always satisfied in our case of interest). In fact, one may probably to adapt the proof of Theorem 7.11 to the case when the graph of groups corresponding to the decomposition of M into pieces is acylindrical. Moreover, such graph of groups is acylindrical whenever M is irreducible (see Lemma 8.1 below).

Also observe that, in our situation, the fundamental groups of the pieces are semihyperbolic in the sense of [AIBr]. Since every free abelian subgroup of a semihyperbolic group is quasi-isometrically embedded, an (apparently difficult) strategy could be to find conditions on a graph of semihyperbolic groups in order to ensure that the fundamental group of the graph is itself semihyperbolic.

Some further discussion of related issues can be found in Section 14.2.

Algebraic properties: irreducible graph manifolds

In this chapter we prove several properties of the fundamental group $\pi_1(M)$ of an irreducible graph manifold M . We first establish the useful fact that the graph of groups decomposition of $\pi_1(M)$ is acylindrical (Section 8.1). This is then used to show that $\pi_1(M)$ is C^* -simple (Section 8.2), that in most cases, $\pi_1(M)$ is SQ-universal (Section 8.3), and that the word problem for $\pi_1(M)$ is always solvable (Section 8.4). Finally, in the last section, we study how the choice of the gluing between pieces can affect $\pi_1(M)$.

8.1. Acylindrical graph of groups

Given a graph of groups the associated fundamental group G has an action on the Bass-Serre tree T . Following Delzant [De], we say the graph of groups is *acylindrical* if it has the property that there exists a constant K , such that the subgroup which pointwise fixes any path in T of length $\geq K$ is automatically trivial. We have the easy:

LEMMA 8.1 (Irreducible \Rightarrow acylindrical). *Let M be an irreducible graph manifold. Then the graph of group description of $\pi_1(M)$ associated to the decomposition of M into pieces is acylindrical.*

PROOF. In fact, one can verify that the graph of group is acylindrical with $K = 3$. Indeed, take any path \mathcal{P} of length three, let e_1, e_2, e_3 be the three consecutive edges in the path \mathcal{P} , and v_1, v_2 the two intermediate vertices. Given a vertex group $G_v = \pi_1(V) \times \mathbb{Z}^k$, we note that there are three possibilities for how an element g in G_v acts on the edges incident to the vertex v , depending on the image of g under the projection map $\rho : G_v \rightarrow \pi_1(V)$:

- if $\rho(g) = e$ (where e is the identity of $\pi_1(M)$), then it fixes every edge adjacent to v ,
- if $\rho(g) \neq e$, and $\rho(g)$ is conjugate to a cusp subgroup in $\pi_1(V)$, then g fixes exactly one incident edge, and permutes all the remaining edges, and
- if $\rho(g) \neq e$, and $\rho(g)$ is **not** conjugate to a cusp subgroup in $\pi_1(V)$, then g can not fix any edge adjacent to v .

To see this trichotomy, recall that $\pi_1(V)$ acts on the universal cover $\widehat{\mathbb{H}}^{n-k}$ of V , which is a copy of hyperbolic space with a suitable $\pi_1(V)$ -equivariant family of (open) horoballs removed. The edges incident to the vertex v in the Bass-Serre tree are in a natural bijective correspondence with the connected components of $\partial(\widehat{\mathbb{H}}^{n-k})$. Moreover, this correspondence is equivariant with respect to the respective actions by G_v . The behavior of an element $g \in G_v$ on the edges incident to v can hence be detected by looking at the action of $\rho(g)$ on the set of connected components of $\partial(\widehat{\mathbb{H}}^{n-k})$. The trichotomy is now obvious.

If we have an element g in $\text{Fix}(\mathcal{P})$, we have that g must fix the two edges e_1, e_2 incident to the vertex v_1 . From the trichotomy above, we conclude that g , when viewed as an element in the stabilizer v_1 , is entirely contained in the \mathbb{Z}^{k_1} factor of the splitting $G_{v_1} = \pi_1(V_1) \times \mathbb{Z}^{k_1}$. Similarly, when viewed as an element in the stabilizer of v_2 , g is contained in the \mathbb{Z}^{k_2} factor of the corresponding splitting $G_{v_2} = \pi_1(V_2) \times \mathbb{Z}^{k_2}$. Now think of g as an element in the stabilizer $G_e \cong \mathbb{Z}^{n-1}$ of the edge e joining v_1 to v_2 . We have that g must lie in the intersection of the two subgroups $\mathbb{Z}^{k_1} \hookrightarrow \mathbb{Z}^{n-1}$ and $\mathbb{Z}^{k_2} \hookrightarrow \mathbb{Z}^{n-1}$. But the intersection of these two subgroups is trivial, since our graph manifold is irreducible. We conclude that g has to be the identity element, as desired. \square

8.2. C^* -simplicity

From the action of the fundamental group G of a graph manifold on the Bass-Serre tree, one can easily deduce a number of additional properties for such groups. As a first application, recall that to any countable discrete group G , one can associate $C_r^*(G)$, its *reduced C^* -algebra*. This algebra is obtained by looking at the action $g \mapsto \lambda_g$ of G on the Hilbert space $l^2(G)$ of square summable complex-valued functions on G , given by the left regular representation:

$$\lambda_g \cdot f(h) = f(g^{-1}h) \quad g, h \in G, \quad f \in l^2(G)$$

The algebra $C_r^*(G)$ is defined to be the operator norm closure of the linear span of the operators λ_g inside the space $B(l^2(G))$ of bounded linear operators on $l^2(G)$. The algebra $C_r^*(G)$ encodes various analytic properties of the group G , and features prominently in the Baum-Connes conjecture. A group G is said to be *C^* -simple* if the algebra $C_r^*(G)$ is a simple algebra, *i.e.* has no proper two-sided ideals. We refer the interested reader to the survey paper by de la Harpe [dlH] for an extensive discussion of this notion.

PROPOSITION 8.2 (Acylindrical $\Rightarrow C^*$ -simple). *Let Γ be the fundamental group of an acylindrical graph of groups, whose Bass-Serre tree has infinitely many ends. Then Γ is C^* -simple.*

PROOF. A criterion for C^* -simplicity was discovered by Powers [Po], who showed that the free group on two generators is C^* -simple. Bridson and de la Harpe abstracted out the geometrical aspects of Power's arguments in [Br-dlH]. Since we will use a criterion from the latter paper, we now introduce some terminology from [Br-dlH]. A homeomorphism g of a Hausdorff space Ω is said to be *dynamically hyperbolic* provided it has exactly two fixed points s, r , and exhibits sink/source dynamics with respect to these fixed points: for any open neighborhood U, V of the points s, r respectively, there exists an N such that $g^n(\Omega \setminus V) \subset U$ and $g^{-n}(\Omega \setminus U) \subset V$ holds for all $n \geq N$. Then [Br-dlH, Corollary 1.2] states that if a group Γ acts by homeomorphisms on a Hausdorff space Ω , and the action satisfies the two properties:

- (i) Γ contains an element g acting hyperbolically on Ω , with fixed points r, s , and
- (ii) there exists a non-trivial element $h \in \Gamma$ with the property that for every integer $i \neq 0$, we have

$$h^i(\text{Stab}(r) \cup \text{Stab}(s))h^{-i} \cap (\text{Stab}(r) \cup \text{Stab}(s)) = \{e\}$$

then the group Γ is C^* -simple.

We now apply the criterion described above in the special case where Γ is the fundamental group of an acylindrical graph of groups. We have a natural action of Γ on the Bass-Serre tree T , and hence an induced action on the boundary at infinity $\Omega := \partial^\infty T$ (a Hausdorff space). This action contains elements which act hyperbolically on the tree T , in the sense that they leave a (unique) geodesic invariant, and act via translations on the invariant geodesic. Let \hat{g} be any such element, and let γ be the corresponding geodesic. The two ends of γ define a pair of points r, s in the boundary at infinity Ω , and it is easy to see that the action of \hat{g} on Ω is dynamically hyperbolic, with fixed points r, s . Since the Bass-Serre tree has infinitely many ends, we can find a second element \hat{h} which acts hyperbolically on the tree T , and has the property that the corresponding invariant geodesic γ' intersects γ in an interval of length $l < \infty$ (if $\gamma \cap \gamma' = \emptyset$, we set $l = -1$).

We now claim that by taking suitable powers of \hat{g} and \hat{h} , we can obtain a pair of elements satisfying conditions (i) and (ii) listed above. More precisely, choose powers n, m so that the translation length of \hat{g}^m and \hat{h}^n are both larger than the integer l , and set $g := \hat{g}^m, h := \hat{h}^n$. Since (i) is immediate, let us focus on establishing (ii). We have the obvious identification:

$$h^i(Stab(r) \cup Stab(s))h^{-i} = Stab(h^i(r)) \cup Stab(h^i(s))$$

so that condition (ii) just states that

$$(8.1) \quad (Stab(h^i(r)) \cup Stab(h^i(s))) \cap (Stab(r) \cup Stab(s))$$

consists solely of the identity element. Moreover, from the way the action on the boundary at infinity is defined, the points $h^i(r), h^i(s)$ correspond precisely to the two ends of the geodesic $h^i(\gamma)$. Of course, if any of the two points $\{h^i(r), h^i(s)\}$ coincides with any of the two points $\{r, s\}$, then the intersection in expression (8.1) would contain the stabilizer of the endpoint of a periodic geodesic, and would hence (at least) contain a \mathbb{Z} . But recall that h was chosen so that the action on γ' was by a translation whose magnitude was *larger* than the length of the intersection $\gamma \cap \gamma'$. This forces the corresponding interval $h^i(\gamma) \cap \gamma'$ to be *disjoint* from the interval $\gamma \cap \gamma'$. In turn, this tells us that the two geodesics γ and $h^i(\gamma)$ are disjoint (by a version of Helly's theorem for trees). This implies that the corresponding endpoints satisfy $\{r, s\} \cap \{h^i(r), h^i(s)\} = \emptyset$, ensuring that the intersection in expression (8.1) actually has a chance of being $\{e\}$.

We now verify that the intersection in expression (8.1) consists solely of the trivial element. This is equivalent to showing that each of the two groups $Stab(r), Stab(s)$ has trivial intersection with each of the two groups $Stab(h^i(r)), Stab(h^i(s))$. Without loss of generality, let us argue that $Stab(r) \cap Stab(h^i(r)) = \{e\}$ (the argument for the remaining three pairs being identical). The pair of points $\{r, h^i(r)\} \subset \Omega$ define a unique geodesic η in the Bass-Serre tree T . An element which fixes both $r, h^i(r)$ must likewise leave η invariant, and moreover must preserve the chosen orientation on η . Note that the intersection $\eta \cap \gamma$ is an infinite geodesic ray. Identifying a geodesic with a copy of \mathbb{R} (equipped with the standard simplicial structure with vertices at \mathbb{Z}), we get a natural homomorphism from the stabilizer of the geodesic to the group of simplicial automorphisms of the real line, $Simp(\mathbb{R})$. The kernel of this homomorphism is the subgroup which fixes the geodesic pointwise, which must be trivial, since we have an acylindrical amalgamation. So the stabilizer of an arbitrary geodesic in T has to be isomorphic to a subgroup of $Simp(\mathbb{R}) \cong \mathbb{D}_\infty$, the

infinite dihedral group. We conclude that stabilizers of geodesics are either trivial, or isomorphic to one of the groups \mathbb{Z}_2 , \mathbb{Z} , or \mathbb{D}_∞ . Moreover, the subgroup which stabilizes the geodesic *and* preserves its orientation is either trivial or isomorphic to \mathbb{Z} (the remaining two groups contain elements that reverse orientation). One can perform a similar analysis for a geodesic ray in T , and see that the semigroup stabilizing the geodesic ray must be either trivial, or isomorphic to the semigroup \mathbb{N} consisting of the natural numbers under addition (as a convention, we include $0 \in \mathbb{N}$).

Let $Stab_o$ denote the subgroup which stabilizes a geodesic *and* preserves the orientation. Applying the analysis from the previous paragraph to the geodesic γ , we have that $Stab_o(\gamma) \cong \mathbb{Z}$. We want to show that $Stab_o(\eta) = Stab(r) \cap Stab(h^i(r))$ is trivial. If it is non-trivial, then from the previous paragraph, we conclude that $Stab_o(\eta) \cong \mathbb{Z}$. The geodesic ray $\eta \cap \gamma$ is stabilized by a sub-semigroup of $Stab_o(\gamma)$ isomorphic to \mathbb{N} , so has to have an infinite stabilizer. From the previous paragraph, we conclude that $Stab(\gamma \cap \eta) \cong \mathbb{N}$. Since $Stab_o(\eta) \cong \mathbb{Z}$ acts on η by translation, it has a sub-semigroup isomorphic to \mathbb{N} contained in $Stab(\gamma \cap \eta)$. We now have two sub-semigroups of $Stab(\gamma \cap \eta)$, given by $Stab_o(\gamma) \cap Stab(\gamma \cap \eta)$ and $Stab_o(\eta) \cap Stab(\gamma \cap \eta)$. But both these sub-semigroups are themselves isomorphic to \mathbb{N} , and hence must have non-trivial intersection inside $Stab(\gamma \cap \eta) \cong \mathbb{N}$. Such a non-trivial element lies in both $Stab_o(\eta)$ and $Stab_o(\gamma)$, so must fix the points at infinity corresponding to both the geodesics. This gives an element acting hyperbolically on the tree T , yet fixing three distinct points r, s , and $h^i(r)$ on the boundary at infinity of the tree, a contradiction. We conclude that $Stab(r) \cap Stab(h^i(r)) = \{e\}$, as desired. The same argument applied to the three remaining pairs of points completes the verification of condition (ii) in Bridson and de la Harpe, giving us that Γ is C^* -simple, as desired. □

Combining Lemma 8.1 and Proposition 8.2, we can show that fundamental groups of irreducible graph manifolds are almost always C^* -simple. More precisely, we have:

COROLLARY 8.3. *If M is an irreducible graph manifold. Then $\pi_1(M)$ is C^* -simple, except in the case where M consists of a single piece, with non-trivial torus fiber, and no internal walls.*

PROOF. If M has at least two pieces, or if it has one piece with an internal wall, then the corresponding Bass-Serre tree has infinitely many ends, and Proposition 8.2 implies $\pi_1(M)$ is C^* -simple. In the remaining case, $\pi_1(M)$ is of the form $\pi_1(V) \times \mathbb{Z}^k$, where V is a non-compact, finite volume hyperbolic manifold with toral cusps. If the torus fiber is trivial, i.e. $k = 0$, then $\pi_1(M) = \pi_1(V)$ is a non-elementary, relatively hyperbolic group. For these groups, Arzhantseva and Minasyan [AM] have shown that being C^* -simple is equivalent to having no non-trivial finite normal subgroup. Since $\pi_1(M)$ is torsion-free, this latter condition is automatically satisfied, and hence $\pi_1(M)$ is indeed C^* -simple.

So we are left with the case where $\pi_1(M) \cong \pi_1(V) \times \mathbb{Z}^k$, and the torus fiber is non-trivial, i.e. $k \geq 1$. Then it is clear that the group $\pi_1(M)$ has non-trivial center. But C^* -simple groups automatically have trivial center, so we conclude that $\pi_1(M)$ is *not* C^* -simple, finishing the proof of the Corollary. □

8.3. SQ-universality

Recall that a group G is *SQ-universal* if every countable group can be embedded into a quotient of G . Olshanskii [Ol] showed that non-elementary Gromov hyperbolic groups are SQ-universal. Again, the simple example of \mathbb{Z}^n shows that CAT(0)-groups are not in general SQ-universal. Our next result is:

PROPOSITION 8.4. *Let M be an irreducible graph manifold, and assume that M does not consist of a single piece with a single internal wall (i.e. the associated graph of groups is distinct from the one vertex one edge graph). Then $\pi_1(M)$ is SQ-universal.*

PROOF. If M consists of a single piece, then $\pi_1(M) \cong \pi_1(V) \times \mathbb{Z}^k$, where V is a finite volume hyperbolic manifold of dimension ≥ 3 . Since $\pi_1(V)$ is a non-elementary (properly) relatively hyperbolic group, work of Arzhantseva, Minasyan and Osin [AMO] implies that $\pi_1(V)$ is SQ-universal. Since $\pi_1(M)$ surjects onto a SQ-universal group, it is itself SQ-universal.

So we may now assume that $\pi_1(M)$ supports a splitting as a non-trivial amalgamation (or an HNN-extension). Lyndon & Schupp [LySc] provide some criterions under which an amalgamation or HNN-extension is SQ-universal. For a group A , define a *blocking set* for a subgroup $C \leq A$ to be a pair of distinct elements $\{x, y\} \subset A \setminus C$ with the property that all the intersections $x^{\pm 1}Cy^{\pm 1} \cap C = \{1\}$. Then [LySc, pg. 289, Theorem V.11.3] establishes that, if the subgroup C is blocked inside A , the amalgamation $G = A *_C B$ is SQ-universal. Similarly, if G is obtained as an HNN extension by identifying two subgroups C_1, C_2 (with $C_i \leq A$), and one can find an element $z \in A$ such that $zC_1z^{-1} \cap C_1 = \{1\} = zC_2z^{-1} \cap C_2$, then the HNN-extension $G = A *_C C_1 = C_2$ is SQ-universal (see [LySc, pg. 292, Theorem V.11.7]).

We now verify that the conditions for SQ-universality are fulfilled for the amalgamations (or HNN extensions) that arise in our situation. Let us first consider the case of an amalgamation: the group $\pi_1(M)$ splits as an amalgamation over $C := \mathbb{Z}^{n-1}$, with the two vertex groups A, B themselves fundamental groups of graph manifolds (with fewer pieces than M). Since the amalgamating subgroup \mathbb{Z}^{n-1} is contained in a piece, it is sufficient to show that a blocking set exists within the fundamental group of that piece. By projecting onto the first factor, the group $\pi_1(V) \times \mathbb{Z}^k$ acts on $\widehat{\mathbb{H}}^{n-k}$, a copy of hyperbolic space with a suitable $\pi_1(V)$ -equivariant collection of (open) horoballs removed. The subgroup $C = \mathbb{Z}^{n-1}$ can then be identified with the subgroup that leaves invariant a fixed boundary horosphere $\mathcal{H} \subset \partial(\widehat{\mathbb{H}}^{n-k})$. In this context, the blocking condition requires us to find two elements $x, y \in \pi_1(V) \setminus C$ with the property that $x^{\pm 1}Cy^{\pm 1} \cap C = \{1\}$, which is equivalent to $(x^{\pm 1}Cy^{\pm 1}) \cdot \mathcal{H} \neq \mathcal{H}$. The $\pi_1(V)$ action on $\widehat{\mathbb{H}}^{n-k}$ is via isometries, so it is sufficient to show that we can find elements x, y having the property that the following sets of distances satisfy:

$$\{d(x^{\pm 1} \cdot \mathcal{H}, \mathcal{H})\} \cap \{d(y^{\pm 1} \cdot \mathcal{H}, \mathcal{H})\} = \emptyset.$$

Now pick $x \in \pi_1(V)$ stabilizing some horosphere \mathcal{H}' (distinct from \mathcal{H}). Then we know that x does not leave any other horosphere invariant, so $d(x^{\pm 1} \cdot \mathcal{H}, \mathcal{H}) > 0$. Moreover, taking large powers of x , we can find an n for which the two real numbers $d(x^{\pm n} \cdot \mathcal{H}, \mathcal{H})$ are as large as we want. In particular, there exists a sufficiently large

$n \in \mathbb{N}$ such that, for $y := x^n$, the distance $d(y^{\pm 1} \cdot \mathcal{H}, \mathcal{H})$ exceeds the distances $d(x^{\pm 1} \cdot \mathcal{H}, \mathcal{H})$.

By the discussion in the previous paragraph, this implies that $\{x, y\}$ form a blocking set for the \mathbb{Z}^{n-k-1} subgroup in $\pi_1(V)$ corresponding to the stabilizer of the horosphere \mathcal{H} . Taking the product with any element in the \mathbb{Z}^k factor gives a blocking set for the subgroup \mathbb{Z}^{n-1} inside $\pi_1(V) \times \mathbb{Z}^k$. This completes the verification of SQ-universality in the case where $\pi_1(M)$ is expressed as an amalgamation.

In the case where $\pi_1(M)$ is obtained purely from HNN extensions (i.e. the underlying graph of groups has *no* separating edge), one can proceed similarly. Pick an edge in the graph of groups. The fact that the edge fails to separate means that, cutting M along the corresponding codimension one torus N , we obtain a connected graph manifold M' . Correspondingly, we can realize $\pi_1(M)$ as an HNN-extension with vertex group $\pi_1(M')$. To check the condition in Lyndon & Schupp, we need to verify that given the two subgroups $C_1, C_2 \leq \pi_1(M')$ corresponding to two boundary components of M' , one can find a $z \in \pi_1(M')$ satisfying $zC_iz^{-1} \cap C_i = \emptyset$. But from our hypothesis, the graph manifold M' is itself non-trivial, i.e. is not of the form $N \times T^k$ for some finite volume hyperbolic manifold N . As a result, $\pi_1(M')$ has a non-trivial decomposition as a graph of groups. In terms of the action of $\pi_1(M')$ on its Bass-Serre tree T , the subgroups C_i lie in the stabilizer of two vertices v_i (we allow the case $v_1 = v_2$), and do *not* stabilize any edge in T . Since the tree T is infinite, we can pick an edge e in T , with the property that both vertices v_i are at distance ≥ 2 from the endpoints of e . Letting $G_e \cong \mathbb{Z}^{n-1}$ be the stabilizer of that edge, we will argue that one can choose an element $z \in G_e$ with the desired property.

If $z \in G_e$ is an arbitrary element, and w_j are the two distinct endpoints of e , the z -action on the edges incident to w_j is determined by the image of z under the projection map $\rho_j : G_{w_j} \cong \pi_1(V_j) \times \mathbb{Z}^{k_j} \rightarrow \pi_1(V_j)$ (see the discussion in the proof of Lemma 8.1). Since $\rho(z)$ is obviously contained in a cusp subgroup, there are only two possibilities:

- $\rho_j(z) = e$, and the z -action fixes every edge incident to w_j , or
- $\rho_j(z) \neq e$, and the z -action permutes all the remaining edges.

The two projection maps ρ_j , when restricted to the group $G_e \cong \mathbb{Z}^{n-1}$ have kernels which are free abelian subgroups of rank $k_j \leq n-3$. The union of these two kernels cannot be the entire group, so we choose an element $z \in G_e$ having non-trivial projection onto the two $\pi_1(V_j)$, and verify that this element has the desired property.

Recall that the subgroups C_i that are being amalgamated lie in the stabilizer of two distinct vertices $v_i \in T$, which are at distance ≥ 2 from both of the vertices w_j . Our element z was chosen so that the z -action maps e to itself, and permutes all the other edges adjacent to e . As a result, the only vertices that are left invariant under z are exactly the two endpoints w_j of the edge e , forcing $z \cdot v_i \neq v_i$. Since $zC_iz^{-1} \leq zG_{v_i}z^{-1} = G_{z \cdot v_i}$, we see that $zC_iz^{-1} \cap C_i \leq G_{z \cdot v_i} \cap G_{v_i}$. As the vertices $v_i, z \cdot v_i$ are distinct, there is a non-trivial path \mathbf{p} joining them, and the subgroup $G_{z \cdot v_i} \cap G_{v_i}$ is contained in the subgroup $Fix(\mathbf{p}) \leq \pi_1(M')$ which fixes the path \mathbf{p} pointwise. The vertices $v_i, z \cdot v_i$ are each at distance ≥ 2 from the edge e . Since z does not fix any edge distinct from e and exiting from an endpoint of e , the path \mathbf{p} must pass through one of the endpoints of e , allowing us to conclude that the

length of \mathbf{p} is ≥ 4 . As M' is an irreducible graph manifold, Lemma 8.1 (see also the first line of its proof) implies that the subgroup $Fix(\mathbf{p})$ is trivial. We conclude that each of the subgroups $zC_i z^{-1} \cap C_i$ is likewise trivial, and hence that the HNN-extension $\pi_1(M) \cong \pi_1(M') *_{C_1=C_2}$ is SQ-universal. This completes the proof of the proposition. \square

REMARK 8.5. We note that, in the proof above, the irreducibility of the graph manifold was only used to deal with the case where the associated graph of groups has no separating edge. As such, the argument given above also shows that if M is an arbitrary (i.e. not necessarily irreducible) graph manifold, *which contains at least one internal wall that separates the manifold into two components*, then the fundamental group $\pi_1(M)$ is SQ-universal

8.4. Solvable word problem

We now shift our attention to an algorithmic question. Given a finite presentation of a group G , the *word problem* asks whether there exists an algorithm for deciding whether or not two words w_1, w_2 in the generators represent the same element in the group G . Building on work of Dehn, who resolved the case where G is a surface group, we know that this question is equivalent to the presentation having a recursive Dehn function (see Gersten [Ge]). In the case where the group G is the fundamental group of a compact connected Riemannian manifold (possibly with boundary), a consequence of the well known Filling Theorem (see e.g. Burillo and Taback [BuTa]) is that the presentation has a recursive Dehn function if and only if the 2-dimensional filling function for the universal cover \widetilde{M} has a recursive upper bound. As we will require this in our arguments, we remind the reader of the definition of the 2-dimensional filling function:

$$Area_M(L) := \sup_c \inf_D \{ Area(D) \mid D : \mathbb{D}^2 \rightarrow \widetilde{M}, D|_{\partial\mathbb{D}^2} = c, L(c) \leq L \} .$$

In other words, we find a minimal area spanning disk for each curve, and try to maximize this area over all curves of length $\leq L$. We are now ready to show:

PROPOSITION 8.6. *Let M be a compact manifold, and assume that M contains an embedded finite family of pairwise disjoint 2-sided smooth submanifolds N_i , cutting M into a finite collection of connected open submanifolds M_j (denote by \bar{M}_j their closure). Moreover, assume this decomposition has the following properties:*

- (a) *each inclusion $N_i \hookrightarrow \bar{M}_j$, and $\bar{M}_j \hookrightarrow M$ is π_1 -injective,*
- (b) *each $\pi_1(N_i)$ is a quasi-isometrically embedded subgroup of $\pi_1(M)$, and*
- (c) *each $\pi_1(M_j)$ has solvable word problem.*

Then the group $\pi_1(M)$ also has solvable word problem.

PROOF. To show that $\pi_1(M)$ has solvable word problem, we need to find a recursive function $F : \mathbb{N} \rightarrow \mathbb{N}$ having the property that, if $\gamma : S^1 \rightarrow \widetilde{M}$ is any closed curve of length $\leq n$, one can find a bounding disk $H : \mathbb{D}^2 \rightarrow \widetilde{M}$ with area $\leq F(n)$. This will be achieved by giving a construction for finding a bounding disk, and verifying that the resulting areas are bounded above by a recursive function.

From hypothesis (a), $\pi_1(M)$ is the fundamental group of a graph of groups \mathcal{G} , with vertex groups isomorphic to the various $\pi_1(M_j)$, and edge groups isomorphic

to the various $\pi_1(N_i)$. Let T denote the associated Bass-Serre tree. Take closed tubular neighborhoods $\widehat{N}_i \supset N_i$ be of the various N_i , chosen small enough so as to be pairwise disjoint. Let $\widehat{M}_j \supset M_j$ be the manifold with boundary obtained by taking the union of \widehat{M}_j with all of the various \widehat{N}_i (ranging over all N_i that occur as boundary components of \widehat{M}_j). The inclusion $M_j \subset \widehat{M}_j$ is clearly a π_1 -isomorphism.

Next, let us construct a map from M to the graph \mathcal{G} . This is achieved by mapping each $\widehat{N}_i \cong N_i \times [-1, 1]$ to the edge labelled by the corresponding $\pi_1(N_i)$, by first collapsing \widehat{N}_i onto the interval factor $[-1, 1]$, and then identifying the interval with the edge. Finally, each connected component of the complement $M \setminus \bigcup \widehat{N}_i$ is entirely contained inside one of the submanifolds M_j ; we map the component to the vertex $v_j \in \mathcal{G}$ whose label is $\pi_1(M_j)$. This map lifts to an equivariant map $\Phi : \widetilde{M} \rightarrow T$, which we will use to analyze the behavior of a closed loop $\gamma : S^1 \rightarrow \widetilde{M}$. Note that Φ is essentially the map defining the “tree of spaces” structure on \widetilde{M} , see Section 2.3 (particularly the discussion around Definition 2.11).

Our analysis of the loop γ will start by associating a *type* to each point in S^1 , *i.e.* by defining a map from S^1 to the vertex set of T . Using the map $\Phi \circ \gamma$, we first assign the type of any point lying in the pre-image of a vertex $v \in T$ to be that same vertex. We now need to discuss how to extend this map to points in the preimage of an open edge $e^\circ \subset T$ (*i.e.* e° excludes the two endpoints of e). Each connected component of the pre-image of e° is either the whole S^1 , or an open interval $U = (a, b)$ in the circle, which inherits an orientation from the ambient S^1 . In the first case, we choose an endpoint v of e , and we simply establish that every point of S^1 has type equal to v . Otherwise, the two endpoints of the interval $U = (a, b)$ either (i) map to the same vertex v in T , or (ii) map to distinct vertices v, w in T . In case (i), we define the type of that interval to be the vertex v . In case (ii), taking into account the orientation on the interval, we can talk of an “initial vertex” $\Phi(\gamma(a)) = v$, and a “terminal vertex” $\Phi(\gamma(b)) = w$. The restriction of γ to $U = (a, b)$ maps into a subset \widehat{N}_i . Let $t \in (a, b)$ be the largest t so that $\gamma(t) \in N_i$. Then we define the type of the points in $(a, t]$ to be v , and the type of the points in (t, b) to be w . By construction, we have that the type function $\rho : S^1 \rightarrow \text{Vert}(T)$ takes on values contained in the image of $\Phi \circ \gamma(S^1)$, and hence only assumes *finitely many* values (as the latter set is compact).

Let us now fix a vertex v of T . Having defined the type function $\rho : S^1 \rightarrow \text{Vert}(T)$ associated to the closed loop γ , we now have that either $\rho^{-1}(v)$ is equal to the whole S^1 , or the preimage $\rho^{-1}(v)$ satisfies the following properties:

- (1) each connected component of $\rho^{-1}(v)$ is a half-open interval $(a_k, b_k] \subset S^1$, and there are finitely many such components,
- (2) there exists a fixed connected lift \widetilde{M}_j of some \widehat{M}_j with the property that the restriction of γ to each connected component $(a_k, b_k]$ has image α_k contained entirely inside \widetilde{M}_j ,
- (3) the point $\gamma(a_k)$ lies on the lift W_k of some N_j , and the point $\gamma(b_k)$ lies on the lift W'_k of some (possibly distinct) $N_{j'}$, and
- (4) if one considers the cyclically ordered collection of intervals $(a_k, b_k]$ along the circle S^1 , then we have that $W'_k = W_{k+1}$.

Except for the fact that there are finitely many components in $\rho^{-1}(v)$ (which will be justified later), the four properties stated above follow immediately from the

definition of the type function ρ . Let us concentrate on the case when $\rho^{-1}(v)$ is not the whole S^1 , the case when ρ is constant being much easier. We proceed to construct a bounding disk for γ , where γ has length $\leq L \in \mathbb{N}$, and to estimate the resulting area. This will be achieved by first expressing γ as a concatenation of loops γ_v , where v ranges over all the (finitely many) types associated to the loop γ . The bounding disk for γ will be obtained by concatenating the bounding disks for the γ_v .

So let $v \in \text{Vert}(T)$ lie in the range of the type function, and consider the connected lift \widetilde{M}_j given by property (2). Each W_k appearing in property (3) is a connected lift of one of the N_i . From hypothesis (a), W_k is a copy of the universal cover of N_i , and from hypothesis (b), the inclusion $W_k \hookrightarrow \widetilde{M}$ is a quasi-isometric embedding. As there are only finitely many such N_i in M , we can choose constants $C, K \in \mathbb{N}$ so that all the inclusions $W_k \hookrightarrow \widetilde{M}$ are (C, K) -quasi-isometries. The two points $\gamma(b_{k-1}) \in W'_{k-1}$ and $\gamma(a_k) \in W_k$ are contained in the same W_k by property (4); let β_k be a minimal length curve in W_k joining them together. The distance between these two points is clearly $\leq L$ in \widetilde{M} , so as measured inside the submanifold W_k , their distance is $\leq CL + K$. Define the loop γ_v by cyclically concatenating $\alpha_1 * \beta_1 * \alpha_2 * \beta_2 * \cdots * \alpha_r * \beta_r$. Since each of the β_i has length $\leq CL + K$, while the union of the α_i has length $\leq L$ (being a subpath of the loop γ), we can estimate the total length of γ_v to be $\leq r \cdot (CL + K) + L \in \mathbb{N}$.

So to complete our estimate on the length of γ_v , we need to estimate the integer r (this will also justify the ‘‘finitely many’’ in property (1) above). For any of the intervals $U = (a_k, b_k] \subset S^1$ in $\rho^{-1}(v)$, the type of the point a_k is a vertex w which is adjacent to v . Correspondingly, there is another subinterval $V \subset S^1$, consisting of points of type w , which satisfies $V \cap \bar{U} = \{a_k\}$. Moreover, there exists a small neighborhood $[a_k - \epsilon, a_k + \delta] \subset V \cup U$ whose image under γ lies entirely in a connected lift \widehat{N}_i of some \widehat{N}_i , and whose endpoints map to opposite boundary components of $\widehat{N}_i \cong \widetilde{N}_i \times [-1, 1]$. For each of the $\widehat{N}_i \subset M$, we let $\lambda_i > 0$ denote the minimal distance between the two boundary components of $\widehat{N}_i \cong N_i \times [-1, 1]$. Since there are only finitely many such \widehat{N}_i , we can find a $\lambda \in \mathbb{N}$ so that $1/\lambda \leq \min\{\lambda_i\}$. We have seen above that to each connected component inside each of the sets $\rho^{-1}(v)$ (where $v \in \text{Vert}(T)$), we can associate a subpath of γ contained inside a connected lift of one of the \widehat{N}_i , which moreover connects opposite boundary components of the lift. These paths are pairwise disjoint, and from the discussion above, has length $\geq 1/\lambda$. We conclude that the total number of such paths is bounded above by $\lambda \cdot L \in \mathbb{N}$. In particular, this gives us the upper bound $\lambda \cdot L$ for:

- the number r of connected components in $\rho^{-1}(v)$, for any $v \in \text{Vert}(T)$, and
- the total number of vertices $v \in \text{Vert}(T)$ for which $\rho^{-1}(v)$ is non-empty.

Combining this with our estimate above, we see that the total length of γ_v is bounded above by the natural number $\lambda CL^2 + \lambda KL + L$.

From hypothesis (a), the space \widetilde{M}_j can be identified with the universal cover of \widehat{M}_j . From hypothesis (c), $\pi_1(M_j)$ has solvable word problem, and hence the 2-dimensional filling function $\text{Area}_{\widehat{M}_j}$ on \widetilde{M}_j has a recursive upper bound $F_j : \mathbb{N} \rightarrow \mathbb{N}$. Observe that there are only finitely many \widehat{M}_j inside the manifold M , hence we can

choose a single recursive $F : \mathbb{N} \rightarrow \mathbb{N}$ which serves as a common upper bound for *all* the 2-dimensional filling functions for the $\widetilde{\widetilde{M}}_j$ (for instance, take $F = \sum F_j$). Then we can find a bounding disk for γ_v whose area is $\leq F(\lambda CL^2 + \lambda KL + L)$. Finding such a bounding disk for each of the vertices v in the range of the type map ρ , we obtain a bounding disk for the original curve γ . As we know that there are $\leq \lambda \cdot L$ vertices in the range of ρ , we conclude that the original curve γ has a bounding disk of total area

$$\leq \lambda \cdot L \cdot F(\lambda CL^2 + \lambda KL + L)$$

Finally, we recall that the class of recursive functions is closed under composition as well as elementary arithmetic operations, and hence the function

$$G(L) := \lambda \cdot L \cdot F(\lambda CL^2 + \lambda KL + L)$$

provides the desired recursive upper bound for the function $Area_M$. From the Filling Theorem [BuTa], we conclude that $\pi_1(M)$ has a recursive Dehn function, and hence that the word problem is solvable for $\pi_1(M)$. □

Note that the obvious decomposition of a graph manifold into pieces satisfies property (a) in the statement of the previous Proposition. Moreover, since all the pieces support a locally CAT(0) metric, their fundamental groups have solvable word problem (see for instance Bridson and Haefliger [BrHa, Section 3.Γ, Theorem 1.4]), so property (c) always holds. Finally, if the graph manifold is assumed to be irreducible, then properties (b) also holds (see Theorem 7.11). This gives us the immediate:

COROLLARY 8.7 (Irreducible \Rightarrow solvable word problem). *For M an irreducible graph manifold, the fundamental group $\pi_1(M)$ has solvable word problem.*

Remark: (1) The above proposition doesn't seem to appear in the literature, though it is no doubt well-known to experts. Indeed, estimates for the Dehn function of a free product with amalgam (or HNN-extension) in terms of the Dehn functions of the vertex groups along with estimates of the *relative distortion* of the edge group inside the vertex groups first seems to have been studied in the (unpublished) thesis of A. Bernasconi [Be]. See also the stronger estimates recently obtained by Arzhantseva and Osin [AO].

(2) The argument given in the proposition shows that, assuming all vertex groups have solvable word problem, the complexity of the word problem for the fundamental group of a graph of groups is closely related to the distortion of the edge/vertex groups in the ambient group (see also the discussion in Farb [Fa1]). In fact, one can weaken hypothesis (c) in the statement of the proposition by instead requiring the distortion of each $\pi_1(N_i)$ inside $\pi_1(M)$ to be bounded above by a recursive function (generalizing the linear bound one has in the special case of a QI-embedding). The same argument works to show that $\pi_1(M)$ still has solvable word problem.

8.5. Gluings and isomorphism type

In this final section, we consider the question of when the fundamental groups of a pair of graph manifolds are isomorphic. Let us first recall that, by Theorem 4.4,

a pair M_1, M_2 of graph manifolds can have isomorphic groups only if there is a bijection between the pieces of M_1 and the pieces of M_2 , having the property that the bijection respects the fundamental groups of the pieces. This implies that the only possible freedom occurs in the *gluing maps*, telling us how the various pieces are glued to each other.

For the sake of simplicity, we will only treat the case when the pieces involved are constructed starting from cusped hyperbolic manifolds of a fixed dimension $n \geq 3$. Let us fix a finite *directed* graph \mathcal{G} , that is a finite connected CW-complex of dimension one with an orientation attached to every edge, and let \mathcal{V}, \mathcal{E} be the sets of vertices and edges of \mathcal{G} . As usual, the valency of a vertex v of \mathcal{G} is the total number of germs of edges starting or ending at v . For each $v \in \mathcal{V}$ with valency h let N_v be a (truncated) cusped hyperbolic n -manifold with at least h cusps. Let $k \leq n - 2$ be fixed, and set $V_v = N_v \times T^k$. We define $G_v = \pi_1(V_v) = \pi_1(N_v) \times \mathbb{Z}^k$, and we associate to every germ of edge starting or ending at v a subgroup $H_{e,v}$ of G_v , in such a way that the following conditions hold:

- each $H_{e,v}$ is (a fixed representative in the conjugacy class of) the fundamental group of a boundary component of V_v ;
- $H_{e,v}$ is not conjugated to $H_{e',v}$ whenever $e \neq e'$, *i.e.* subgroups corresponding to different edges with an endpoint in v are associated to different boundary components of V_v .

As a consequence, every $H_{e,v}$ is isomorphic to \mathbb{Z}^{n+k-1} . The graph \mathcal{G} and the groups $G_v, H_{e,v}$ determine what we call a *pregraph* of groups.

For every $e \in E$ let now $v_-(e), v_+(e) \in \mathcal{V}$ be respectively the starting point and the ending point of e . A *gluing pattern* for \mathcal{G} is a collection of group isomorphisms $\Phi = \{\varphi_e: H_{e,v_-(e)} \rightarrow H_{e,v_+(e)}, e \in \mathcal{E}\}$. We say that Φ is irreducible if for every $e \in \mathcal{E}$ the fiber subgroup of $H_{e,v_+(e)}$ intersects trivially the image of the fiber subgroup of $H_{e,v_-(e)}$ via φ_e . Of course, every gluing pattern for \mathcal{G} defines a graph of groups (\mathcal{G}, Φ) , which has in turn a well-defined fundamental group $\pi_1(\mathcal{G}, \Phi)$, according to the Bass-Serre theory. We say that (\mathcal{G}, Φ) is supported by \mathcal{G} , and is irreducible if Φ is.

Let $\mathcal{M}(\mathcal{G})$ be the set of diffeomorphism classes of graph manifolds obtained by gluing the pieces $V_v, v \in \mathcal{V}$ according to the pairing of the boundary components encoded by the edges of \mathcal{G} . It follows by Theorem 0.5 that the isomorphism classes of fundamental groups of (irreducible) graph of groups supported by \mathcal{G} coincide with the isomorphism classes of fundamental groups of (irreducible) manifolds in $\mathcal{M}(\mathcal{G})$.

REMARK 8.8. The assumption $k \leq n - 2$ on the dimensions of toric and hyperbolic factors of the pieces will play a crucial role in the proof of Theorem 8.9 below. Note however that there could not exist irreducible gluing patterns for \mathcal{G} if the dimension of the toric factors of the pieces exceeded the dimension of the hyperbolic factors. Moreover, it seems reasonable (and the proof of Theorem 8.9 strongly suggests) that an analogue of Theorem 8.9 could also hold when different pieces have toric factors of variable dimensions, provided that such dimensions are sufficiently small.

The main result of this section is the following:

THEOREM 8.9. *Suppose that \mathcal{G} has at least two vertices. Then, there exist infinitely many irreducible graphs of groups supported by \mathcal{G} with mutually non-isomorphic fundamental groups. Equivalently, there exist infinitely many diffeomorphism classes of irreducible manifolds in $\mathcal{M}(\mathcal{G})$.*

PROOF. An automorphism of a pregraph of groups is a combinatorial automorphism φ of \mathcal{G} (as an *undirected* graph) such that $G_{\varphi(v)}$ is isomorphic to G_v for every $v \in \mathcal{V}$ (as discussed at the beginning of the proof of Lemma 5.1, this is equivalent to requiring that $V_{\varphi(v)}$ is diffeomorphic to V_v for every $v \in \mathcal{V}$). We say that a pre-graph of groups is *without symmetries* if it does not admit non-trivial automorphisms. We first consider the case when \mathcal{G} is without symmetries.

Since \mathcal{G} has at least two vertices, there exists an edge $e \in \mathcal{E}$ with distinct endpoints $v_1 = v_-(e)$, $v_2 = v_+(e)$. We fix this edge for use in the rest of the proof.

Let Φ , Φ' be irreducible gluing patterns for \mathcal{G} . Consider $\varphi: H_{e,v_1} \rightarrow H_{e,v_2}$ (resp. $\varphi': H_{e,v_1} \rightarrow H_{e,v_2}$) the isomorphism of Φ (resp. of Φ') associated to the edge e . We say that Φ' is *equivalent* to Φ if there exist an automorphism ψ_1 of G_{v_1} and an automorphism ψ_2 of G_{v_2} such that $\psi_1(H_{e,v_1}) = H_{e,v_1}$, $\psi_2(H_{e,v_2}) = H_{e,v_2}$ and $\varphi' \circ \psi_1|_{H_{e,v_1}} = \psi_2|_{H_{e,v_2}} \circ \varphi$. Note that this notion of equivalence is only sensitive to the behavior of the gluing along the single edge e , and completely ignores what happens along the remaining edges in \mathcal{G} .

Now, the proof of Theorem 8.9 will follow immediately from the following two facts:

Fact 1: If $\pi_1(\mathcal{G}, \Phi) \cong \pi_1(\mathcal{G}, \Phi')$, then Φ is equivalent to Φ' .

Fact 2: There exist infinitely many pairwise non-equivalent irreducible gluing patterns for \mathcal{G} .

Let us begin by establishing **Fact 1**. Let $\psi: \pi_1(\mathcal{G}, \Phi) \rightarrow \pi_1(\mathcal{G}, \Phi')$ be a group isomorphism. By Theorem 4.4, the isomorphism ψ induces an automorphism of \mathcal{G} . But by hypothesis, we are in the case where \mathcal{G} has no symmetries, so the automorphism of \mathcal{G} must be the identity. In particular, we have $\psi(G_1) = g_1 G'_1 g_1^{-1}$, $\psi(G_2) = g_2 G'_2 g_2^{-1}$, where G_i (resp. G'_i) is the image of G_{v_i} in $\pi_1(\mathcal{G}, \Phi)$ (resp. in $\pi_1(\mathcal{G}, \Phi')$), and g_1, g_2 are elements in $\pi_1(\mathcal{G}, \Phi')$. If H (resp. H') is the image in $\pi_1(\mathcal{G}, \Phi)$ (resp. in $\pi_1(\mathcal{G}, \Phi')$) of H_{e,v_1} and H_{e,v_2} (which are identified by the very definition of fundamental group of a graph of groups), since ψ induces the identity of \mathcal{G} we also have $\psi(H) = g_3 H' g_3^{-1}$ for some $g_3 \in \pi_1(\mathcal{G}, \Phi')$.

Up to conjugating ψ , we can assume $g_1 = 1$, so that $\psi(G_1) = G'_1$. Next note that we have $g_3 H' g_3^{-1} = \psi(H) \subseteq \psi(G_1) = G'_1$, so $H' \subseteq g_3^{-1} G'_1 g_3 \cap G'_1$. By Lemma 2.14-(5), this implies that either $g_3 \in G'_1$, or H' corresponds to an edge of \mathcal{G} having both endpoints on the vertex representing G'_1 . But recall that the edge e was chosen to have distinct endpoints, ruling out this last possibility. So at the cost of conjugating ψ with g_3^{-1} , we may further assume that $g_3 = 1$, and both $\psi(G_1) = G'_1$ and $\psi(H) = H'$. As a consequence we have $H' = \psi(H) \subseteq \psi(G_2) = g_2 G'_2 g_2^{-1}$, so $H' \subseteq g_2 G'_2 g_2^{-1} \cap G_2$, whence $g_2 \in G'_2$ as above and $\psi(G_2) = G'_2$.

We have thus proved that ψ induces isomorphisms $G_1 \cong G'_1$, $G_2 \cong G'_2$ which “agree” on $H = G_1 \cap G_2$. More precisely, for $i = 1, 2$ there exists an isomorphism $\psi_i: G_{v_i} \rightarrow G_{v_i}$ such that the following conditions hold: $\psi_i(H_{e,v_i}) = H_{e,v_i}$ for $i = 1, 2$, and $\varphi' \circ \psi_1|_{H_{e,v_1}} = \psi_2|_{H_{e,v_2}} \circ \varphi$. By definition, this means that Φ is equivalent to Φ' , and **Fact 1** is proved.

Let us now prove **Fact 2**. Recall that for $i = 1, 2$ we have an identification $G_{v_i} \cong \Gamma_i \times \mathbb{Z}^k$, where $\Gamma_i = \pi_1(N_{v_i})$. We also denote by L_i the subgroup of Γ_i such that $L_i \times \mathbb{Z}^k < \Gamma_i \times \mathbb{Z}^k$ corresponds to H_{e, v_i} under the above identification. As showed in the proof of Lemma 5.1, every automorphism of $G_{v_i} = \Gamma_i \times \mathbb{Z}^k$ is of the form $(g, v) \mapsto (\theta_i(g), \alpha_i(g) + \beta_i(v))$, where $\theta_i: \Gamma_i \rightarrow \Gamma_i$ and $\beta_i: \mathbb{Z}^k \rightarrow \mathbb{Z}^k$ are isomorphisms, and $\alpha_i: \Gamma_i \rightarrow \mathbb{Z}^k$ is a homomorphism. We now claim that, in a sense to be made precise below, if we restrict to automorphisms leaving L_i invariant, then there exist at most a finite number of possibilities for the isomorphism θ_i .

Let $\tilde{\Theta}_i$ be the group of automorphisms of Γ_i leaving L_i invariant, and let Θ_i be the group of automorphisms of L_i given by restrictions of elements of $\tilde{\Theta}_i$. For $g \in \Gamma_i$, we denote by $c_g \in \text{Aut}(\Gamma_i)$ the conjugation by g . If $\theta, \hat{\theta} \in \tilde{\Theta}_i$ are such that $\theta = c_g \circ \hat{\theta}$ for some $g \in \Gamma_i$, then $gL_i g^{-1} = L_i$, whence $g \in L_i$ (see the proof of Lemma 2.14-(1)). Since L_i is abelian, this implies that θ and $\hat{\theta}$ restrict to the same element of Θ_i . As a consequence, Θ_i has at most the cardinality of the group of outer automorphisms of Γ_i , which is finite by Mostow rigidity (together with the well-known fact that the group of isometries of a complete finite-volume hyperbolic manifold is finite). We have thus proved the fact claimed above that Θ_i is finite.

For $i = 1, 2$, let us now fix a free basis of $L_i \times \mathbb{Z}^k \cong \mathbb{Z}^{n+k-1}$ whose first $n-1$ elements give a basis of L_i and whose last k elements give a basis of \mathbb{Z}^k . Under the induced identification of L_i with \mathbb{Z}^{n-1} , the group Θ_i is identified with a finite subgroup of $\text{SL}(n-1, \mathbb{Z})$, which will still be denoted by Θ_i . Moreover, we may identify the group of automorphisms of $H_{e, v_i} \cong L_i \times \mathbb{Z}^k$ with the group of matrices $\text{SL}(n+k-1, \mathbb{Z})$. The discussion above shows that under these identifications every automorphism of $L_i \times \mathbb{Z}^k$ which extends to an automorphism of G_{v_i} has the form

$$\left(\begin{array}{c|c} \theta_i & 0 \\ \hline v_i & w_i \end{array} \right) \in \text{SL}(n+k-1, \mathbb{Z}), \quad \theta_i \in \Theta_i < \text{SL}(n-1, \mathbb{Z}),$$

and any isomorphism between $\varphi: H_{e, v_1} \rightarrow H_{e, v_2}$ may be represented by a matrix

$$\left(\begin{array}{c|c} A & B \\ \hline C & D \end{array} \right) \in \text{SL}(n+k-1, \mathbb{Z}),$$

where A, D have order $(n-1) \times (n-1)$ and $k \times k$ respectively. Moreover, it is readily seen that φ can be extended to an irreducible gluing pattern if and only if $\text{rk}(B) = k$.

Now, since $k < n-1$ and Θ_2 is finite, it is possible to construct an infinite sequence $\{B_j\}_{j \in \mathbb{N}}$ of matrices of order $(n-1) \times k$ such that the following conditions hold:

- $\text{rk} B_j = k$ for every $j \in \mathbb{N}$;
- if Λ_j is the subgroup of \mathbb{Z}^{n-1} generated by the columns of B_j , $j \in \mathbb{N}$, and $\Lambda_j = \theta(\Lambda_h)$ for some $\theta \in \Theta_2$, then necessarily $j = h$.

Let $\varphi_j: H_{e, v_1} \rightarrow H_{e, v_2}$, $j \in \mathbb{N}$, be the isomorphism represented by the matrix

$$P_j = \left(\begin{array}{c|c} \text{Id}_{n-1} & B_j \\ \hline 0 & \text{Id}_k \end{array} \right),$$

and extend φ_j to an irreducible gluing pattern Φ_j . We now claim that Φ_j is not equivalent to Φ_h if $j \neq h$, thus concluding the proof of (2). In fact, if Φ_j is equivalent

to Φ_h , then there exist matrices

$$N_1 = \left(\begin{array}{c|c} \theta_1 & 0 \\ \hline v_1 & w_1 \end{array} \right), \quad N_2 = \left(\begin{array}{c|c} \theta_2 & 0 \\ \hline v_2 & w_2 \end{array} \right)$$

such that $\theta_i \in \Theta_i$, $w_i \in \mathrm{SL}(k, \mathbb{Z})$ for $i = 1, 2$, and $P_j N_1 = N_2 P_h$. It is readily seen that this condition implies the equality $B_j w_1 = \theta_2 B_h$. Since $w_1 \in \mathrm{SL}(k, \mathbb{Z})$, this implies in turn $\Lambda_j = \theta_2(\Lambda_h)$, whence $j = h$ by the properties of the B_j 's listed above. We have thus proved the theorem under the assumption that \mathcal{G} is without symmetries.

In the general case, the arguments just described ensure that an infinite family $\{\Phi_i\}_{i \in \mathbb{N}}$ of irreducible gluing patterns exists such that, if $i \neq j$, then $\pi_1(\mathcal{G}, \Phi_i)$ is not isomorphic to $\pi_1(\mathcal{G}, \Phi_j)$ via an isomorphism inducing the identity of \mathcal{G} . Suppose now by contradiction that the groups $\pi_1(\mathcal{G}, \Phi_i)$ fall into finitely many isomorphism classes. Then, up to passing to an infinite subfamily, we may suppose that for every $i, j \in \mathbb{N}$ there exists an isomorphism $\psi_{ij}: \pi_1(\mathcal{G}, \Phi_i) \rightarrow \pi_1(\mathcal{G}, \Phi_j)$ inducing the automorphism δ_{ij} of \mathcal{G} . Since the group of automorphisms of \mathcal{G} is finite, there exist $h, k \in \mathbb{N} \setminus \{0\}$ such that $h \neq k$ and $\delta_{0h} = \delta_{0k}$. Therefore, the map $\psi_{0k} \circ \psi_{0h}^{-1}$ establishes an isomorphism between $\pi_1(\mathcal{G}, \Phi_h)$ and $\pi_1(\mathcal{G}, \Phi_k)$ inducing the identity of \mathcal{G} , a contradiction. \square

REMARK 8.10. The assumption that \mathcal{G} has at least two vertices is not really necessary. In the case that \mathcal{G} has only one vertex, we could provide a different proof of Theorem 8.9 just by replacing our analysis of isomorphisms between amalgamated products with an analogous analysis of isomorphisms between HNN-extensions.

REMARK 8.11. The strategy described in the proof of Theorem 8.9 can also be applied to the examples discussed in Remark 12.11, where an infinite family $\{M_i\}_{i \geq 1}$ of irreducible manifolds not supporting any CAT(0) metric is constructed by gluing two fixed 4-dimensional pieces V_1, V_2 along their unique boundary component. With notation as in Corollary 12.10 and Remark 12.11, we now show that if V_1 is not diffeomorphic to V_2 , then M_i is not diffeomorphic to M_j for every $i, j \in \mathbb{N}$, $i \neq j$.

Let us choose bases for the fundamental groups of the boundary components of V_1, V_2 (such components are 3-dimensional tori) in such a way that the first vector is null-homologous in V_i , $i = 1, 2$, and the last one belongs to the fiber subgroup (which is isomorphic to \mathbb{Z}). Then the gluing map defining M_n is encoded by the matrix

$$A_n = \begin{pmatrix} 1 & * & 1 \\ 0 & * & 0 \\ 0 & * & n \end{pmatrix}.$$

Moreover, every homomorphism of the fundamental group of a piece into the fiber subgroup (which is abelian) vanishes on null-homologous elements, whence on horizontal slopes. So any automorphism of the fundamental group of each of the two pieces, when restricted to the boundary, gives an automorphism of the form

$$\begin{pmatrix} * & * & 0 \\ * & * & 0 \\ 0 & * & \pm 1 \end{pmatrix}$$

(see the proof of Theorem 8.9). It is now readily seen that if N_1, N_2 are matrices of this form, then we have

$$N_1 A_n = \begin{pmatrix} * & * & * \\ * & * & * \\ * & * & \pm n \end{pmatrix} \neq \begin{pmatrix} * & * & * \\ * & * & * \\ * & * & \pm m \end{pmatrix} = A_m N_2.$$

Now, since V_1 is not diffeomorphic to V_2 , the M_i 's are associated to a graph without symmetries. As explained in the proof Theorem 8.9, this is now sufficient to conclude that the M_i 's are pairwise non-diffeomorphic.

Also, observe that by the proof of Theorem 8.9, if V_1 is diffeomorphic to V_2 we can still conclude that among the M_i 's there exist infinitely many pairwise non-diffeomorphic manifolds.

Pieces of irreducible graph manifolds are quasi-preserved

In this chapter, we prove Theorem 0.12, which we recall here for the convenience of the reader:

THEOREM (Pieces of irreducible manifolds are preserved). *Let M_1, M_2 be a pair of irreducible graph manifolds, and $\Gamma_i = \pi_1(M_i)$ their respective fundamental groups. Let $\Lambda_1 \leq \Gamma_1$ be a subgroup conjugate to the fundamental group of a piece in M_1 , and $\varphi : \Gamma_1 \rightarrow \Gamma_2$ be a quasi-isometry. Then, the set $\varphi(\Lambda_1)$ is within finite Hausdorff distance from a conjugate of $\Lambda_2 \leq \Gamma_2$, where Λ_2 is the fundamental group of a piece in M_2 .*

So, let us fix graph manifolds M_1, M_2 with fundamental groups $\Gamma_i = \pi_1(M_i)$ and suppose $\psi : \Gamma_1 \rightarrow \Gamma_2$ is a quasi-isometry. Due to Milnor-Svarc Lemma (see Theorem 1.2), ψ induces a quasi-isometry between \widetilde{M}_1 and \widetilde{M}_2 , which we will still denote by ψ . The statement of Theorem 0.12 is equivalent to the fact that ψ sends, up to a finite distance, chambers of \widetilde{M}_1 into chambers of \widetilde{M}_2 . In order to prove this fact, we will use the technology of *asymptotic cones*, which we now briefly describe.

9.1. The asymptotic cone of a geodesic metric space

Roughly speaking, the asymptotic cone of a metric space gives a picture of the metric space as “seen from infinitely far away”. It was introduced by Gromov in [Gr1], and formally defined in [vdDWi].

A *filter* on \mathbb{N} is a set $\omega \subseteq \mathcal{P}(\mathbb{N})$ satisfying the following conditions:

- (1) $\emptyset \notin \omega$;
- (2) $A, B \in \omega \implies A \cap B \in \omega$;
- (3) $A \in \omega, B \supseteq A \implies B \in \omega$.

For example, the set of complements of finite subsets of \mathbb{N} is a filter on \mathbb{N} , known as the *Fréchet filter* on \mathbb{N} .

A filter ω is a *ultrafilter* if for every $A \subseteq \mathbb{N}$ we have either $A \in \omega$ or $A^c \in \omega$, where $A^c := \mathbb{N} \setminus A$. For example, fixing an element $a \in \mathbb{N}$, we can take the associated *principal ultrafilter* to consist of all subsets of \mathbb{N} which contain a . An ultrafilter is *non-principal* if it does not contain any finite subset of \mathbb{N} .

It is readily seen that a filter is an ultrafilter if and only if it is maximal with respect to inclusion. Moreover, an easy application of Zorn’s Lemma shows that any filter is contained in a maximal one. Thus, non-principal ultrafilters exist (just take any maximal filter containing the Fréchet filter).

From this point on, let us fix a non-principal ultrafilter ω on \mathbb{N} . As usual, we say that a statement \mathcal{P}_n depending on $n \in \mathbb{N}$ holds ω -a.e. if the set of indices such that \mathcal{P}_n holds belongs to ω . If X is a topological space, and $(x_n) \subseteq X$ is a sequence

in X , we say that $\omega\text{-lim } x_n = x_\infty$ if $x_n \in U$ ω -a.e. for every neighbourhood U of x_∞ . When X is Hausdorff, an ω -limit of a sequence, if it exists, is unique. Moreover, any sequence in any compact space admits a ω -limit. For example, any sequence (a_n) in $[0, +\infty]$ admits a unique ω -limit.

Now let (X_i, x_i, d_i) , $i \in \mathbb{N}$, be a sequence of pointed metric spaces. Let \mathcal{C} be the set of sequences $(y_n), y_n \in X_n$, such that $\omega\text{-lim } d_n(x_n, y_n) < +\infty$, and consider the following equivalence relation on \mathcal{C} :

$$(y_n) \sim (z_n) \iff \omega\text{-lim } d_n(y_n, z_n) = 0.$$

We set $\omega\text{-lim}(X_i, x_i, d_i) = \mathcal{C}/\sim$, and we endow $\omega\text{-lim}(X_i, x_i, d_i)$ with the well-defined distance given by $d_\omega([(y_n)], [(z_n)]) = \omega\text{-lim } d_n(y_n, z_n)$. The pointed metric space $(\omega\text{-lim}(X_i, x_i, d_i), d_\omega)$ is called the ω -limit of the pointed metric spaces X_i .

Let (X, d) be a metric space, $(x_n) \subseteq X$ a sequence of base-points, and $(r_n) \subset \mathbb{R}^+$ a sequence of rescaling factors diverging to infinity. We introduce the notation $(X_\omega((x_n), (r_n)), d_\omega) := \omega\text{-lim}(X_n, x_n, d/r_n)$.

DEFINITION 9.1. The metric space $(X_\omega((x_n), (r_n)), d_\omega)$ is the *asymptotic cone* of X with respect to the ultrafilter ω , the basepoints (x_n) and the rescaling factors (r_n) . For conciseness, we will occasionally just write $X_\omega((x_n), (r_n))$ for the asymptotic cone, the distance being implicitly understood to be d_ω .

If ω is fixed and $(a_n) \subseteq \mathbb{R}$ is any sequence, we say that (a_n) is $o(r_n)$ (resp. $O(r_n)$) if $\omega\text{-lim } a_n/r_n = 0$ (resp. $\omega\text{-lim } |a_n|/r_n < \infty$).

Let $(x_n) \subseteq X$, $(r_n) \subseteq \mathbb{R}$ be fixed sequences of basepoints and rescaling factors, and set $X_\omega = (X_\omega((x_n), (r_n)), d_\omega)$. Sequences of subsets in X give rise to subsets of X_ω : if for every $n \in \mathbb{N}$ we are given a subset $\emptyset \neq A_n \subseteq X$, we set

$$\omega\text{-lim } A_n = \{[(p_n)] \in X_\omega \mid p_n \in A_n \text{ for every } n \in \mathbb{N}\}.$$

It is easily seen that for any choice of the A_n 's, the set $\omega\text{-lim } A_n$ is closed in X_ω . Moreover, $\omega\text{-lim } A_n \neq \emptyset$ if and only if $\omega\text{-lim } \frac{d(x_n, A_n)}{r_n} < \infty$.

9.2. Quasi-isometries and asymptotic cones

We are interested in describing how quasi-isometries asymptotically define bi-Lipschitz homeomorphisms. In order to do this, and to fix some notations, we recall some basic results about ω -limits of quasi-isometries and quasi-geodesics.

Suppose that (Y_i, y_i, d_i) , $i \in \mathbb{N}$ are pointed metric spaces, and that (X, d) is a metric space. Let $(x_n) \subseteq X$ be a sequence of basepoints and $(r_n) \subset \mathbb{R}$ a sequence of rescaling factors. Until the end of the section, to simplify the notation, we set $X_\omega := (X_\omega((x_n), (r_n)))$. The following result is well-known (and very easy):

LEMMA 9.2. *Suppose $(k_i) \subseteq \mathbb{R}^+$, $(c_i) \subseteq \mathbb{R}^+$ are sequences satisfying $k = \omega\text{-lim } k_i < \infty$, and $c_i = o(r_i)$. For each $i \in \mathbb{N}$, let $f_i: Y_i \rightarrow X$ be a map with the property that for every $y, y' \in Y_i$, the inequality*

$$d(f_i(y), f_i(y')) \leq k_i d_i(y, y') + c_i$$

holds. If $d(f_i(y_i), x_i) = O(r_i)$, then the formula $[(p_i)] \mapsto [f_i(p_i)]$ provides a well-defined map $f_\omega: \omega\text{-lim}(Y_i, y_i, d_i/r_i) \rightarrow X_\omega$. Moreover, f_ω is k -Lipschitz, whence continuous. If $k > 0$ and

$$d(f_i(y), f_i(y')) \geq \frac{d_i(y, y')}{k_i} - c_i$$

is also satisfied (i.e. if f_i is a (k_i, c_i) -quasi-isometric embedding), then f_ω is a k -bi-Lipschitz embedding.

As a corollary, quasi-isometric metric spaces have bi-Lipschitz homeomorphic asymptotic cones. We recall that a (k, c) -quasi-geodesic in X is a (k, c) -quasi-isometric embedding of a (possibly unbounded) interval in X .

LEMMA 9.3. *Suppose $(k_i) \subseteq \mathbb{R}^+$, $(c_i) \subseteq \mathbb{R}^+$ are sequences satisfying $k = \omega\text{-lim } k_i < \infty$, and $c_i = o(r_i)$. For each $i \in \mathbb{N}$, let $\gamma_i: [a_i, b_i] \rightarrow X$ be a (k_i, c_i) -quasi-geodesic with image $H_i = \text{Im } \gamma_i$, and assume that $d(x_i, H_i) = O(r_i)$. Then up to precomposing γ_i with a translation of \mathbb{R} , we may suppose that 0 is the basepoint of $[a_i, b_i]$, and that the sequence (γ_i) induces a k -bi-Lipschitz path*

$$\gamma_\omega: [\omega\text{-lim}(a_i/r_i), \omega\text{-lim}(b_i/r_i)] \rightarrow X_\omega.$$

Moreover, we have $\text{Im } \gamma_\omega = \omega\text{-lim } H_i$.

PROOF. The only non-trivial (but easy) assertion is the last one, which we leave to the reader. \square

The previous Lemma asserts that quasi-geodesics in a space give rise to a bi-Lipschitz path in the asymptotic cone. Our next result shows a type of converse to this result.

LEMMA 9.4. *Assume X is a geodesic space, and let $\gamma_\omega: [0, 1] \rightarrow X_\omega$ be a k -Lipschitz path. Then, for every $\varepsilon > 0$ there exists a sequence of $(k + \varepsilon)$ -Lipschitz paths $\gamma_n: [0, r_n] \rightarrow X$ with the following properties:*

- $d(x_n, \gamma_n(0)) = O(r_n)$, so if $[0, 1] = \omega\text{-lim } \frac{1}{r_n}[0, r_n]$ then (γ_n) defines a $(k + \varepsilon)$ -Lipschitz path $\omega\text{-lim } \gamma_n: [0, 1] \rightarrow X_\omega$;
- $\omega\text{-lim } \gamma_n = \gamma_\omega$.

PROOF. For every $t \in [0, 1]$ set $p^t = \gamma_\omega(t) = [(p_i^t)]$, and for every $j \in \mathbb{N}$ let $A_j \subseteq \mathbb{N}$ be the set of indices $i \in \mathbb{N}$ such that

$$d(p_i^t, p_i^{t'}) \leq \left(1 + \frac{\varepsilon}{k}\right) r_i d_\omega(p^t, p^{t'})$$

for every $t = h \cdot 2^{-j}, t' = h' \cdot 2^{-j}$, $h, h' \in \mathbb{Z}$, $0 \leq h, h' \leq 2^j$. By construction we have $A_{j+1} \subseteq A_j$ and $A_j \in \omega$ for every $j \in \mathbb{N}$. For every $i \in \mathbb{N}$, let

$$j(i) = \sup\{j \in \mathbb{N} \mid i \in A_j\} \in \mathbb{N} \cup \{\infty\},$$

and set $j'(i) = i$ if $j(i) = \infty$ and $j'(i) = j(i)$ otherwise. By the nature of the construction, we have $i \in A_{j'(i)}$. For every $i \in \mathbb{N}$, we define the curve $\gamma_i: [0, r_i] \rightarrow X$ as follows: if $h \in \{0, 1, \dots, 2^{j'(i)} - 1\}$, then the restriction of γ_i to the interval $[hr_i 2^{-j'(i)}, (h+1)r_i 2^{-j'(i)}]$ is a linear parameterization of a geodesic joining $p_i^{h 2^{-j'(i)}}$ with $p_i^{(h+1) 2^{-j'(i)}}$. Since $i \in A_{j'(i)}$ each such restriction is $(k + \varepsilon)$ -Lipschitz, so γ_i is $(k + \varepsilon)$ -Lipschitz. It readily follows that $\omega\text{-lim } \gamma_i$ is $(k + \varepsilon)$ -Lipschitz, and in particular continuous. Thus, in order to show that $\omega\text{-lim } \gamma_n = \gamma_\omega$ it is sufficient to show that $(\omega\text{-lim } \gamma_n)(t) = \gamma_\omega(t) = p^t$ for every t of the form $h 2^{-j}$, $h, j \in \mathbb{N}$. However, if $t = h 2^{-j}$ by construction we have

$$\begin{aligned} \{i \in \mathbb{N} \mid \gamma_i(t) = p_i^t\} &\supseteq \{i \mid j \leq j(i) < \infty\} \cup (\{i \mid j(i) = \infty\} \cap \{i \mid i \geq j\}) \\ &\supseteq \{i \mid j \leq j(i)\} \cap \{i \mid i \geq j\} \\ &\supseteq A_j \cap \{i \mid i \geq j\} \in \omega. \end{aligned}$$

As a result, for each $t = h2^{-j}$, we have that $\omega\text{-lim } \gamma_i(t) = [(\gamma_i(t))] = [(p_i^t)] = \gamma_\omega(t)$, whence the conclusion. \square

9.3. Tree-graded spaces

We are going to need some results about the asymptotic cones of complete hyperbolic manifolds of finite volume. The following definitions are taken from [DrSa].

DEFINITION 9.5. A geodesic metric space X is said to be *tree-graded* with respect to a collection of closed subsets $\{P_i\}_{i \in I}$, called *pieces*, if

- (1) $\bigcup P_i = X$,
- (2) $|P_i \cap P_j| \leq 1$ if $i \neq j$,
- (3) if Δ is a simple geodesic triangle contained in X , then it is contained in a single piece.

DEFINITION 9.6. A geodesic metric space X is *asymptotically tree-graded* with respect to a collection of subsets $A = \{H_i\}_{i \in I}$ if the following conditions hold:

- (1) for each choice of basepoints $(x_n) \subseteq X$ and rescaling factors (r_n) , the associated asymptotic cone $X_\omega = X_\omega((x_n), (r_n))$ is tree-graded with respect to the collection of subsets $\mathcal{P} = \{\omega\text{-lim } H_{i(n)} \mid H_{i(n)} \in A\}$, and
- (2) if $\omega\text{-lim } H_{i(n)} = \omega\text{-lim } H_{j(n)}$, where $i(n), j(n) \in I$, then $i(n) = j(n)$ ω -a.e.

We summarize in the following lemma some properties of tree-graded spaces which are proved in [DrSa] and will be useful later.

LEMMA 9.7. *Let P, P' be distinct pieces of a tree-graded space Y . Then there exist $p \in P, p' \in P'$ such that the following holds: for any continuous path $\gamma: [0, 1] \rightarrow Y$ with $\gamma(0) \in P$ and $\gamma(1) \in P'$ we have $p, p' \in \text{Im } \gamma$. Moreover, if $P \cap P' \neq \emptyset$, then $\{p\} = \{p'\} = P \cap P'$. In particular, if there exists $\gamma: [0, 1] \rightarrow Y$ with $\gamma(0) \in P, \gamma(1) \in P'$ and $P \cap \text{Im } \gamma = \{\gamma(0)\}$, then every path joining P and P' passes through $\gamma(0)$.*

The following result is a consequence of [Fa2] and [DrSa] (see also Section 9.6 below):

THEOREM 9.8. *Let B be a neutered space obtained as the complement in \mathbb{H}^n of an equivariant family K of pairwise disjoint open horoballs, and let H be the collection of the boundaries of the horoballs in K . Then B , endowed with its path metric, is asymptotically tree-graded with respect to H .*

\square

Another proof of this theorem (for spaces of pinched negative curvature and without the equivariance assumption) can be found in [Si].

9.4. Studying the asymptotic cones of \widetilde{M}

Let M be an *irreducible* graph manifold with universal covering \widetilde{M} . Let ω be any non-principal ultrafilter on \mathbb{N} , let $(x_n) \subseteq \widetilde{M}, (d_n) \subseteq \mathbb{R}$ be fixed sequences of basepoints and rescaling factors, and set $\widetilde{M}_\omega = (\widetilde{M}_\omega, (x_n), (d_n))$.

DEFINITION 9.9. An ω -chamber (resp. ω -wall, ω -fiber) in \widetilde{M}_ω is a subset $X_\omega \subseteq \widetilde{M}_\omega$ of the form $X_\omega = \omega\text{-lim } X^i$, where each $X^i \subseteq \widetilde{M}$ is a chamber (resp. a wall, a fiber).

Since a constant k exists such that each point of \widetilde{M} has distance at most k from some wall, every point of \widetilde{M}_ω lies in some ω -wall. We will show that every bi-Lipschitz homeomorphism of \widetilde{M}_ω preserves the decomposition of \widetilde{M}_ω into ω -walls. Together with an argument which allows us to recover quasi-isometries of the original spaces from homeomorphisms of asymptotic cones, this will allow us to prove Theorem 0.12.

The following lemma is a direct consequence of Corollary 7.13.

LEMMA 9.10. *There exists $k \geq 1$ such that every ω -wall of \widetilde{M} is k -bi-Lipschitz homeomorphic to \mathbb{R}^{n-1} , and every ω -fiber of \widetilde{M} is k -bi-Lipschitz homeomorphic to \mathbb{R}^h , $h \leq n - 3$.*

From Corollary 7.13, Lemma 9.7 and Theorem 9.8 we get the following:

LEMMA 9.11. *There exists $k \geq 1$ such that for any ω -chamber $C_\omega = \omega\text{-lim } C_i \subseteq \widetilde{M}_\omega$ there exists a k -bi-Lipschitz homeomorphism $\varphi: C_\omega \rightarrow Y \times \mathbb{R}^l$, where Y is a tree-graded space whose pieces are k -bi-Lipschitz homeomorphic to \mathbb{R}^{n-l-1} . Moreover, suppose that W_i, W'_i are walls adjacent to C_i , and let $W_\omega = \omega\text{-lim } W_i \neq \emptyset$, $W'_\omega = \omega\text{-lim } W'_i \neq \emptyset$. Then, there exist pieces P_ω, P'_ω of Y such that $W_\omega = \varphi^{-1}(\pi^{-1}(P_\omega))$ and $W'_\omega = \varphi^{-1}(\pi^{-1}(P'_\omega))$, where $\pi: Y \times \mathbb{R}^l \rightarrow Y$ is the natural projection. Therefore:*

- if $W_\omega = W'_\omega$ then $W_i = W'_i$ ω -a.e.
- if $W_\omega \neq W'_\omega$ then there exists a fiber $F_\omega \subseteq W_\omega$ such that every continuous path joining a point in W_ω with a point in W'_ω has to pass through F_ω . Moreover, $W_\omega \cap W'_\omega$ is either empty or equal to F_ω .

LEMMA 9.12. *Let $W_\omega, W'_\omega \subseteq \widetilde{M}_\omega$ be distinct ω -walls. Then there exists an ω -fiber $F_\omega \subseteq W_\omega$ with the following properties:*

- (1) every Lipschitz path joining a point in W_ω with a point in W'_ω passes through F_ω ;
- (2) $W_\omega \cap W'_\omega \subseteq F_\omega$.

PROOF. Let $\gamma: [0, 1] \rightarrow \widetilde{M}_\omega$ be a Lipschitz path with $\gamma(0) \in W_\omega$, $\gamma(1) \in W'_\omega$, and let $W_i, W'_i \subseteq \widetilde{M}$, $i \in \mathbb{N}$, be walls such that $\omega\text{-lim } W_i = W_\omega$, $\omega\text{-lim } W'_i = W'_\omega$. Since $W_\omega \neq W'_\omega$, we may suppose $W_i \neq W'_i$ for every $i \in \mathbb{N}$.

Let us take $\varepsilon > 0$. By Lemma 9.4, $\gamma = \omega\text{-lim } \gamma_i$ where $\gamma_i: [0, a_i] \rightarrow \widetilde{M}$ is a $(k + \varepsilon)$ -Lipschitz path and $\omega\text{-lim } a_i/d_i < \infty$. Of course (see the proof of Lemma 9.4) we may suppose $\gamma_i(0) \in W_i$, $\gamma_i(a_i) \in W'_i$ ω -a.e.

For every $i \in \mathbb{N}$, let us define a wall L_i and a chamber C_i as follows: if both W_i and W'_i are adjacent to the same chamber, then $L_i = W'_i$ and C_i is the chamber adjacent both to W_i and to L_i ; if W_i, W'_i do not intersect the same chamber, then $L_i \neq W_i$ and C_i are such that $W_i \cap C_i \neq \emptyset$, $L_i \cap C_i \neq \emptyset$, and every path connecting W_i and W'_i intersects L_i (the existence of such L_i, C_i is an obvious consequence of the realization of \widetilde{M} as a tree of spaces). We would like to associate to γ_i a continuous path α_i joining W_i with L_i which does not intersect any chamber different from C_i . This can be done in the following way. Let z_i be the last point of γ_i which lies on W_i , let p_i be the first point of γ_i which follows z_i and lies on L_i and call γ'_i the subpath of γ_i with endpoints z_i, p_i . We have that $\gamma'_i \cap C_i$ is a collection of paths in C_i , and, since γ'_i is rectifiable and the distance between walls is bounded from below, only finitely many of them, say $\delta_1^i, \dots, \delta_m^i$, have endpoints in different walls.

By concatenating the δ_j^i 's with suitable geodesics ψ_j^i contained in the appropriate thin walls we obtain the desired α_i . By construction, α_i intersects W_i only in its initial point. Also note that because thin walls are quasi-isometrically embedded in \widetilde{M} the length of α_i is uniformly linearly bounded by the length of γ_i' , whence of γ_i . Therefore, we can suppose that there exists $k' > 0$ such that α_i is defined on the same interval as γ_i , and α_i is k' -Lipschitz ω -a.e.

Now consider $C_\omega = \omega\text{-lim } C_i$ and $L_\omega = \omega\text{-lim } L_i$. We find ourselves in the context of Lemma 9.11, which implies that there exists an ω -fiber $F_\omega \subseteq W_\omega$ with the property that every path joining W_ω and L_ω passes through F_ω . Now, by Lemma 9.2, $\alpha = \omega\text{-lim } \alpha_i$ is a continuous path joining W_ω and L_ω , so α necessarily passes through F_ω . Since F_ω only depends on W_i, W_i' , in order to prove (1) it is sufficient to show that γ must also pass through F_ω .

Choose the points $q_i \in \text{Im } \alpha_i$ so that the corresponding $q = [(q_i)] \in \text{Im } \alpha$ is the first point along α which belongs to F_ω . By the definition of γ_i and ψ_j^i , at least one of the following possibilities must hold:

- (i) $q_i \in \gamma_i$ ω -a.e.
- (ii) $q_i \in \psi_{j(i)}^i$ ω -a.e. and $l_{B_i}(\psi_{j(i)}^i|_{q_i}) = o(d_i)$, where $\psi_{j(i)}^i|_{q_i}$ denotes the initial subpath of $\psi_{j(i)}^i$ ending in q_i and l_{B_i} denotes the length of the projection of such a path on the base of C_i ,
- (iii) $q_i \in \psi_{j(i)}^i$ ω -a.e. and $\omega\text{-lim } l_{B_i}(\psi_{j(i)}^i|_{q_i})/d_i > 0$.

In cases (i) and (ii), it is clear that there is a point on $\gamma \cap F_\omega$. So let us now prove that case (iii) cannot occur. Indeed, the sequence of the starting points of the $\psi_{j(i)}^i$'s gives a point $q' \neq q$ which comes before q along α . Since α_i intersects W_i only in its initial point, by Lemma 9.11 the initial subpath of α ending in q' joins a point on W_ω with a point on an ω -wall Q_ω such that $Q_\omega \neq W_\omega$, and, by our hypothesis on q , it does not pass through the fiber F_ω . But the portion of $\omega\text{-lim } \psi_{j(i)}^i$ between q' and q provides a path starting on Q_ω and intersecting W_ω only in $q \in F_\omega$. By Lemma 9.11, this implies that *every* continuous path joining a point on Q_ω to a point in W_ω has to intersect F_ω , a contradiction. This completes the proof of statement (1) in the Lemma.

In order to prove (2), take $p \in W_\omega \cap W'_\omega$. From the definitions, we can express p as $p = [(p_i)] = [(p'_i)]$, where $p_i \in W_i, p'_i \in W'_i$ ω -a.e. and $d(p_i, p'_i) = o(r_i)$. Let γ_i be a geodesic joining p_i with p'_i . Then ω -a.e. the image of γ_i intersects L_i . So, if $q_i \in \text{Im } \gamma_i \cap L_i$ we have $d(p_i, q_i) = o(r_i)$, whence $p = [(q_i)]$. This shows that p belongs to $F_\omega = W_\omega \cap L_\omega$, concluding the proof of the Lemma. \square

PROPOSITION 9.13. *Suppose $f: \mathbb{R}^{n-1} \rightarrow \widetilde{M}_\omega$ is a bi-Lipschitz embedding. Then $f(\mathbb{R}^{n-1})$ is an ω -wall of \widetilde{M}_ω .*

PROOF. Set $A = f(\mathbb{R}^{n-1}) \subseteq \widetilde{M}_\omega$. It is sufficient to show that $A \subseteq W_\omega$ for some ω -wall W_ω : if this is the case, A is open (by invariance of domain) and closed (since f is bi-Lipschitz, whence proper) in the connected set W_ω , so would have to coincide with W_ω .

We first show that if W_ω, W'_ω are distinct ω -walls, then either $W_\omega \cap A \subseteq W'_\omega \cap A$ or $W'_\omega \cap A \subseteq W_\omega \cap A$. In fact, suppose on the contrary that there exist $p, q \in \mathbb{R}^{n-1}$ such that $f(p) \in W_\omega \setminus W'_\omega, f(q) \in W'_\omega \setminus W_\omega$. By Lemma 9.12, an ω -fiber F_ω exists such that any bi-Lipschitz path joining $f(p)$ and $f(q)$ passes through F_ω . Therefore,

$\mathbb{R}^{n-1} \setminus f^{-1}(F_\omega)$ is path disconnected. However, $f^{-1}(F_\omega)$ is bi-Lipschitz equivalent to a subset of \mathbb{R}^h , $h \leq n-3$, whence a contradiction (see, for example, [Do]).

We denote by \mathcal{W} the set of ω -walls of \widetilde{M}_ω , and introduce the transitive relation \preceq on \mathcal{W} defined by $W_\omega \preceq W'_\omega$ if and only if $A \cap W_\omega \subseteq A \cap W'_\omega$. The argument above shows that for every pair of ω -walls W_ω, W'_ω we have either $W_\omega \preceq W'_\omega$, or $W'_\omega \preceq W_\omega$, or possibly both.

Now suppose that (\mathcal{W}, \preceq) admits a countable increasing cofinal set, *i.e.* a family $\{W_\omega^i\}_{i \in \mathbb{N}} \subseteq \mathcal{W}$ satisfying $W_\omega^i \preceq W_\omega^{i+1}$, and with the property that for every $W_\omega \in \mathcal{W}$ there exists $i \in \mathbb{N}$ with $W_\omega \preceq W_\omega^i$. Since \widetilde{M}_ω is the union of its walls, we can then express A as the union $A = \bigcup_{i \in \mathbb{N}} A \cap W_\omega^i$. If there exists i_0 such that $A \cap W_\omega^i = A \cap W_\omega^{i_0}$ for every $i \geq i_0$, then we obtain that $A \subseteq W_\omega^{i_0}$, and we are done. After passing to a subsequence, we may therefore assume that the W_ω^i 's are pairwise distinct. Then Lemma 9.12 implies that, for every $i \in \mathbb{N}$, a ω -fiber $F_\omega^i \subseteq W_\omega^i$ exists such that $A \cap W_\omega^i \subseteq A \cap W_\omega^i \cap W_\omega^{i+1} \subseteq A \cap F_\omega^i$. So we can rewrite A as the union $A = \bigcup_{i \in \mathbb{N}} A \cap F_\omega^i$. On the other hand, being bi-Lipschitz equivalent to a subset of \mathbb{R}^{n-3} , each of the sets $A \cap F_\omega^i$ is a closed subspace of A with empty interior. Together with Baire's Lemma, this finally gives a contradiction, thus concluding the proof in the case where (\mathcal{W}, \preceq) admits a countable increasing cofinal set.

Finally, suppose that (\mathcal{W}, \preceq) does *not* admit any countable increasing cofinal set. Let $P = \{p_i\}_{i \in \mathbb{N}}$ be a dense subset of A . Since for every pair W_ω, W'_ω we have either $W_\omega \preceq W'_\omega$ or $W'_\omega \preceq W_\omega$ (or both), one can construct an increasing sequence $\{W_\omega^i\}_{i \in \mathbb{N}} \subseteq \mathcal{W}$ such that $W_\omega^i \supseteq \{p_0, \dots, p_i\}$ for every $i \in \mathbb{N}$. Since $\{W_\omega^i\}_{i \in \mathbb{N}}$ is *not* cofinal, there must exist an ω -wall $H_\omega \in \mathcal{W}$ such that $A \cap H_\omega \supseteq A \cap W_\omega^i$ for every $i \in \mathbb{N}$. We then have $P \subseteq H_\omega$, forcing $A \subseteq H_\omega$ since H_ω is closed in \widetilde{M}_ω , and we are done. \square

COROLLARY 9.14. *For each k, c , there exists $\beta \geq 0$ (only depending on k, c and the geometry of \widetilde{M}) such that the image of \mathbb{R}^{n-1} under a (k, c) -quasi-isometric embedding in \widetilde{M} is contained in the β -neighborhood of a wall.*

PROOF. By contradiction, take a sequence of (k, c) -quasi-isometric embeddings $f_m : \mathbb{R}^{n-1} \rightarrow \widetilde{M}$ such that for each $m \in \mathbb{N}$ and wall $W \subseteq \widetilde{M}$ we have $f_m(\mathbb{R}^{n-1}) \not\subseteq N_m(W)$, where $N_m(W)$ is the m -neighbourhood of W . Fix a point $p \in \mathbb{R}^{n-1}$. The f_m 's induce a bi-Lipschitz embedding f from the asymptotic cone $\mathbb{R}^{n-1} = \mathbb{R}_\omega^{n-1}((p), (m))$ to the asymptotic cone $\widetilde{M}_\omega(f_m(p), (m))$. (Recall that, if X is a metric space, we denote by $X_\omega((x_m), (r_m))$ the asymptotic cone of X associated to the sequence of basepoints (x_m) and the sequence of rescaling factors (r_m) .) By the previous proposition, there is an ω -wall $W_\omega = \omega\text{-lim } W_m$ such that $f(\mathbb{R}^{n-1}) = W_\omega$. By hypothesis, for each m there is a point $p_m \in \mathbb{R}^{n-1}$ with $d(f_m(p_m), W_m) \geq m$. Set $r_m = d(p_m, p)$. By choosing p_m as close to p as possible, we may assume that no point q such that $d(p, q) \leq r_m - 1$ satisfies $d(f_m(q), W_m) \geq m$, so

$$(9.1) \quad d(f_m(q), W_m) \leq m + k + c \quad \text{for every } q \in \mathbb{R}^{n-1} \text{ s.t. } d(p, q) \leq r_m.$$

Notice that $\omega\text{-lim } r_m/m = \infty$, for otherwise $[(p_m)]$ should belong to $\mathbb{R}_\omega^{n-1}((p), (m))$, $[f_m(p_m)]$ should belong to $\widetilde{M}_\omega((f_m(p)), (m))$, and, since $f(\mathbb{R}^{n-1}) = W_\omega$, we would have $d(f_m(p_m), W_m) = o(m)$.

Let us now change basepoints, and consider instead the pair of asymptotic cones $\mathbb{R}_\omega^{n-1}((p_m), (m))$ and $\widetilde{M}_\omega((f_m(p_m)), (m))$. The sequence (f_m) induces a bi-Lipschitz embedding f' between these asymptotic cones (note that $f \neq f'$, simply

because due to the change of basepoints, f and f' are defined on different spaces with values in different spaces!). Let $A_m = \{q \in \mathbb{R}^{n-1} \mid d(q, p) \leq r_m\}$ and $A_\omega = \omega\text{-lim } A_m \subseteq \mathbb{R}_\omega^{n-1}((p_m), (m))$. Since $\omega\text{-lim } r_m/m = \infty$, it is easy to see that A_ω is bi-Lipschitz homeomorphic to a half-space in \mathbb{R}^{n-1} . Moreover, by (9.1) each point in $f'(A_\omega)$ is at a distance at most 1 from $W'_\omega = \omega\text{-lim } W_i$ (as before, observe that the sets W_ω and W'_ω live in different spaces). Again by Proposition 9.13 we have that $f'(A_\omega) \subseteq f'(\mathbb{R}_\omega^{n-1}((p_m), (m))) = W''_\omega$ for some ω -wall W''_ω . Moreover, since $[(f_m(p_m))] \in W''_\omega \setminus W'_\omega$, we have $W'_\omega \neq W''_\omega$.

By Lemma 9.12 there exists a fiber $F_\omega \subseteq W'_\omega \cap W''_\omega$ such that every path joining a point in W''_ω with a point in W'_ω has to pass through F_ω . Now, if $a \in f'(A_\omega)$ we have $d(a, W'_\omega) \leq 1$, so there exists a geodesic of length at most one joining $a \in W''_\omega$ with some point in W'_ω . Such a geodesic must pass through F_ω , so every point of $f'(A_\omega)$ must be at a distance at most 1 from F_ω . If $h: f'(A_\omega) \rightarrow F_\omega$ is such that $d(b, h(b)) \leq 1$ for every $b \in f'(A_\omega)$, then h is a $(1, 2)$ -quasi-isometric embedding. Therefore the map $g = h \circ f': A_\omega \rightarrow F_\omega$ is a quasi-isometric embedding. But this is not possible, since if $n - 1 > l$ there are no quasi-isometric embeddings from a half space in \mathbb{R}^{n-1} to \mathbb{R}^l (as, taking asymptotic cones, such an embedding would provide an injective continuous function from an open set in \mathbb{R}^{n-1} to \mathbb{R}^l). This completes the proof of the corollary. \square

9.5. Walls and chambers are quasi-preserved by quasi-isometries

We are now ready to conclude the proof of Theorem 0.12. We come back to our original situation, *i.e.* we take irreducible graph n -manifolds M_1, M_2 and we suppose that $f: \widetilde{M}_1 \rightarrow \widetilde{M}_2$ is a given (k, c) -quasi-isometry. We will say that a constant is *universal* if it only depends on k, c and on the geometry of M_1, M_2 . We begin by recalling the following well-known result (see *e.g.* [KaLe3, Corollary 2.6]):

LEMMA 9.15. *Let $f: \mathbb{R}^{n-1} \rightarrow \mathbb{R}^{n-1}$ be an (a, b) -quasi-isometric embedding. Then f is an (a', b') -quasi-isometry, where a', b' only depend on a, b .*

PROPOSITION 9.16. *A universal constant λ exists such that for every wall $W_1 \subseteq \widetilde{M}_1$, there exists a wall $W_2 \subseteq \widetilde{M}_2$ with the property that the Hausdorff distance between $f(W_1)$ and W_2 is $\leq \lambda$. Moreover, W_2 is the unique wall in \widetilde{M}_2 at finite Hausdorff distance from $f(W_1)$.*

PROOF. Since M_1 is irreducible, there exists a (k', c') -quasi-isometry $i: \mathbb{R}^{n-1} \rightarrow W_1$ (where k', c' only depend on the geometry of M_1), and Corollary 9.14 (applied to the quasi-isometric embedding $f \circ i$) ensures that $f(W_1)$ is contained in the β -neighbourhood of W_2 for some wall W_2 , where β is universal. For every $y \in f(W_1)$ let $p(y) \in W_2$ be a point such that $d(y, p(y)) \leq \beta$. It follows easily from Lemma 9.15 that the map $p \circ f|_{W_1}: W_1 \rightarrow W_2$ is a (k'', c'') -quasi-isometry, where k'', c'' are universal. This in turn implies that W_2 is contained in the β' -neighbourhood of $f(W_1)$, where β' is universal. The first statement follows, with $\lambda = \max\{\beta, \beta'\}$. The uniqueness of W_2 is an immediate consequence of Lemma 2.17. \square

Putting together Propositions 9.16 and 4.3 we now get the following result, which concludes the proof of Theorem 0.12:

PROPOSITION 9.17. *There exists a universal constant H such that for every chamber $C_1 \subseteq \widetilde{M}_1$ there exists a unique chamber $C_2 \subseteq \widetilde{M}_2$ such that the Hausdorff*

distance between $f(C_1)$ and C_2 is bounded by H . Moreover, if W_1 is a wall adjacent to C_1 then $f(W_1)$ lies at finite Hausdorff distance from a wall W_2 adjacent to C_2 .

9.6. Thickness and relative hyperbolicity

Our next application concerns the notion of relative hyperbolicity, which we briefly recall. Let us fix a finitely generated group G , and a collection of subgroups H_1, \dots, H_n . There are several equivalent definitions of the notion of relative hyperbolicity of G with respect to H_1, \dots, H_n (see *e.g.* [Fa2], [Gr2], [Os]). Since we are going to describe obstructions to relative hyperbolicity coming from the study of asymptotic cones, we recall the characterization of relative hyperbolicity provided by the following result:

THEOREM 9.18 ([DrSa]). *Let $C(G)$ be any Cayley graph of G , and let us consider each H_i as a subset of $C(G)$. Then, the group G is relatively hyperbolic with respect to H_1, \dots, H_n if and only if $C(G)$ is asymptotically tree-graded (see Definition 9.6) with respect to H_1, \dots, H_n . In this case, we say that each H_i is a peripheral subgroup of G .*

The motivating example of a relatively hyperbolic group is the fundamental group of a non-compact, finite volume, Riemannian manifold with sectional curvature bounded above by some negative constant δ . Such a group will be relatively hyperbolic, with respect to the collection of cusp subgroups (we used this fact in the proof of our Proposition 7.4). The class of relatively hyperbolic groups has been extensively studied by a variety of authors, so a reasonable question is whether the fundamental groups of our irreducible graph manifolds are relatively hyperbolic. The following result is an immediate consequence of Dahmani's Combination Theorem [Da] (see also [Si, Theorem 0.2]):

PROPOSITION 9.19. *Assume the graph manifold M has at least one piece with trivial torus fiber. Then $\pi_1(M)$ is relatively hyperbolic with respect to a finite family of proper subgroups.*

So we are left with considering the case where all pieces of M have non-trivial toric factor. Let us moreover assume that the irreducible graph manifold M has associated graph of groups with at least two vertices. Let us sketch the proof that \widehat{M} , and so any Cayley graph of $\pi_1(M)$, is thick of order 1.

The notion of thickness was introduced by Behrstock, Druţu and Mosher in [BDM] as an obstruction for a metric space to be asymptotically tree-graded, and hence, for a group to be relatively hyperbolic. The simplest such obstruction is being unstricted, *i.e.* having no cut-points in any asymptotic cone (by definition, a metric space is thick of order 0 if it is unstricted). It is readily seen that the product of two unbounded geodesic metric spaces is unstricted. Notable thick metric spaces and groups which are not unstricted include the mapping class group and Teichmüller space (equipped with the Weil-Petersson metric) of most surfaces (see [BDM]), fundamental groups of classical 3-dimensional graph manifolds (see again [BDM]), and the group $\text{Out}(F_n)$ for $n \geq 3$ (see Algom-Kfir [A-K]).

Let us briefly describe what it means for a metric space X to be thick of order at most 1 with respect to a collection of subsets \mathcal{L} . First of all, the family \mathcal{L} is required to "fill" X , that is there must exist a positive constant τ such that the

union of the sets in \mathcal{L} is τ -dense in X (property (N1)). Secondly, a certain coarse connectivity property (denoted by (N2)) must be satisfied: for each $L, L' \in \mathcal{L}$ we can find elements $L_0 = L, L_1, \dots, L_n = L'$ of \mathcal{L} such that $N_\tau(L_i) \cap N_\tau(L_{i+1})$ has infinite diameter, where the constant τ is independent of L, L' . The space X is said to be a τ -network with respect to the family of subspaces \mathcal{L} if conditions (N1) and (N2) hold (with respect to the constant τ). For X to be thick of order at most 1, we need X to be a τ -network with respect a family \mathcal{L} , where each $L \in \mathcal{L}$ is unconstricted (actually the stricter condition that the family \mathcal{L} is uniformly unconstricted is required to hold).

Notice that property (N2) fails if X is asymptotically tree-graded with respect to \mathcal{L} as in that case there are uniform bounds on the diameter of $N_k(L) \cap N_k(L')$ for $L, L' \in \mathcal{L}$ with $L \neq L'$.

PROPOSITION 9.20. *Let M be a graph manifold, with at least one internal wall, and with the property that all pieces have non-trivial fibers. Then \widetilde{M} and $\pi_1(M)$ are both thick of order 1.*

PROOF. Let us first argue that \widetilde{M} is thick of order ≤ 1 . We show that \widetilde{M} is a τ -network with respect to the collection \mathcal{H} of its chambers (for τ large enough). In fact, every point in \widetilde{M} is clearly uniformly close to a chamber (property (N1)). Furthermore, if τ is large enough, then the intersection of two adjacent chambers contains a wall. As walls have infinite diameter, we easily obtain property (N2) as well.

To complete the proof that \widetilde{M} is thick of order ≤ 1 we are only left with proving that \mathcal{H} is uniformly unconstricted. This is true because there exists a uniform constant $k \geq 1$ such that any ω -chamber is k -biLispchitz homeomorphic to the product of a geodesic metric space and some \mathbb{R}^n , $n > 0$.

Finally, we note that, by a result of Drutu, Mozes, and Sapir [DMS, Theorem 4.1], any group which supports an acylindrical action on a tree has the property that every asymptotic cone has a cut point. In view of Lemma 8.1, we conclude that $\pi_1(M)$ has cut points in every asymptotic cones, hence cannot be thick of order 0. This concludes the proof of the Lemma. \square

Therefore, if we assume that every piece of our irreducible graph manifold M has non-trivial torus factor, then $\pi_1(M)$ is either thick of order 0 (when M consists of a single piece without internal walls), or thick of order 1 (when M has at least one internal wall). By [BDM, Corollary 7.9], this has the following nice consequence:

COROLLARY 9.21. *Let M be a graph manifold, with the property that all pieces have non-trivial fibers. Then $\pi_1(M)$ is not hyperbolic relative to any collection of proper subgroups.*

Quasi isometry rigidity, I

This chapter is devoted to the proof of Theorem 0.13. We recall the statement for the convenience of the reader:

THEOREM. *Let N be a complete finite-volume hyperbolic m -manifold, $m \geq 3$, and let Γ be a finitely generated group quasi-isometric to $\pi_1(N) \times \mathbb{Z}^d$, $d \geq 0$. Then there exist a finite-index subgroup Γ' of Γ , a finite-sheeted covering N' of N , a group Δ and a finite group F such that the following short exact sequences hold:*

$$1 \longrightarrow \mathbb{Z}^d \xrightarrow{j} \Gamma' \longrightarrow \Delta \longrightarrow 1,$$

$$1 \longrightarrow F \longrightarrow \Delta \longrightarrow \pi_1(N') \longrightarrow 1.$$

Moreover, $j(\mathbb{Z}^d)$ is contained in the center of Γ' . In other words, Γ' is a central extension by \mathbb{Z}^d of a finite extension of $\pi_1(N')$.

In what follows we will give a proof of Theorem 0.13 under the additional assumption that the cusps of N are toric. However, the attentive reader will observe that all the results needed in the proofs below also hold in the case where N is not assumed to have toric cusps, provided that the walls of the universal covering of $\overline{N} \times T^d$ are quasi-isometrically embedded in the universal covering $B \times \mathbb{R}^d$, where B is the neutered space covering \overline{N} . But this last fact is obvious, since the boundary components of $B \times \mathbb{R}^d$ are totally geodesic (in the metric sense).

So, let us consider the graph manifold with boundary $M = \overline{N} \times T^d$, and observe that Γ is quasi-isometric to $\pi_1(M)$. Moreover, M is obviously irreducible, and the universal covering \widetilde{M} of M is isometric to the Riemannian product $B \times \mathbb{R}^d$, where B is a neutered space. The walls of \widetilde{M} coincide with the boundary components of \widetilde{M} .

10.1. The quasi-action of Γ on \widetilde{M}

As discussed in Section 1.3, a quasi-isometry between Γ and $\pi_1(M)$ induces a k -cobounded k -quasi-action h of Γ on \widetilde{M} for some $k \geq 1$. From this point on, we will fix such a quasi-action. Henceforth, for every $\gamma \in \Gamma$, we will abuse notation, and also denote by γ the corresponding quasi-isometry $h(\gamma): \widetilde{M} \rightarrow \widetilde{M}$.

We want to prove that every quasi-isometry $\gamma: \widetilde{M} \rightarrow \widetilde{M}$, $\gamma \in \Gamma$ can be coarsely projected on B to obtain a quasi-isometry of B . We say that a constant is universal if it depends only on k, H and the geometry of B , where H is such that for every $\gamma \in \Gamma$ and every wall $W \subseteq \widetilde{M}$, the set $\gamma(W)$ is at Hausdorff distance at most H from a wall of \widetilde{M} (see Proposition 9.16).

LEMMA 10.1. *There exists a universal constant H' such that, for each fiber $F = \{b\} \times \mathbb{R}^d \subseteq \widetilde{M}$ and each $\gamma \in \Gamma$, the set $\gamma(F)$ is at Hausdorff distance bounded by H' from a fiber $\overline{F} = \{\overline{b}\} \times \mathbb{R}^d \subseteq \widetilde{M}$.*

PROOF. Let $K \subseteq \overline{N}$ be the cut-locus of \overline{N} relative to $\partial\overline{N}$, i.e. the set of points of \overline{N} whose distance from $\partial\overline{N}$ is realized by at least two distinct geodesics, and let $R' = 2 \sup\{d_{\overline{N}}(p, q) \mid p \in K, q \in \partial\overline{N}\}$. Since \overline{N} is compact, R' is finite, and it is easily seen that for each $p \in \overline{N}$ there exist (at least) two distinct components of $\partial\overline{N}$ whose distance from p is at most R' . This implies that for each fiber F there exist two walls W, W' such that $F \subseteq A_{R'}(W, W') = \{x \in \widetilde{M} \mid d(x, W) \leq R', d(x, W') \leq R'\}$.

Moreover, if O, O' are disjoint horospheres in ∂B , it is easy to see that the diameter of the set $\{b \in B \mid d(b, O) \leq R', d(b, O') \leq R'\}$ is bounded by a constant which only depends on R' . As a consequence, if F is a fiber contained in $A_{R'}(W, W')$ then there exists a universal constant D such that $A_{R'}(W, W') \subseteq N_D(F)$. As quasi-isometries almost preserve walls, there exist a universal constant $R'' \geq R'$ and walls $\overline{W}, \overline{W}'$ such that $\gamma(A_{R'}(W, W')) \subseteq A_{R''}(\overline{W}, \overline{W}')$. It follows that γ restricts to a (k', k') -quasi-isometric embedding of F into $A_{R''}(\overline{W}, \overline{W}')$, where k' is a universal constant. But both F and $A_{R''}(\overline{W}, \overline{W}')$ are quasi-isometric to \mathbb{R}^d , so by Lemma 9.15 the restriction of γ to F defines a quasi-isometry (with universal constants) between F and $A_{R''}(\overline{W}, \overline{W}')$, and this forces the Hausdorff distance between $\gamma(F)$ and a fiber in $A_{R''}(\overline{W}, \overline{W}')$ to be bounded by a universal H' . \square

The above Lemma can be used to define a quasi-action of Γ on B . Recall that \widetilde{M} is isometric to $B \times \mathbb{R}^d$, and fix $\gamma \in \Gamma$. We define a map $\psi(\gamma): B \rightarrow B$ by setting $\psi(\gamma)(b) = \pi_B(\gamma((b, 0)))$ for every $b \in B$, where $\pi_B: \widetilde{M} \cong B \times \mathbb{R}^d \rightarrow B$ is the natural projection, and for $(b, f), (b', f') \in B \times \mathbb{R}^d \cong \widetilde{M}$ we denote by $d_B((b, f), (b', f'))$ the distance in B between b and b' (see Section 2). With a slight abuse of notation, we also denote by d_B the distance on B .

We now show that every $\psi(\gamma)$ is a quasi-isometry (with universal constants). Let $b, b' \in B$ and set $F = \{b\} \times \mathbb{R}^d$ and $F' = \{b'\} \times \mathbb{R}^d$. The Hausdorff distance between $\gamma(F)$ and $\gamma(F')$ is bounded from below by $d_B(b, b')/k - k$, so if $\overline{F}, \overline{F}'$ are fibers with Hausdorff distance bounded by H' from $\gamma(F), \gamma(F')$ respectively, then the Hausdorff distance between \overline{F} and \overline{F}' is at least $d_B(b, b')/k - k - 2H'$. We have therefore

$$\begin{aligned} d_B(\psi(\gamma)(b), \psi(\gamma)(b')) & \geq d_B(\pi_B(\overline{F}), \pi_B(\overline{F}')) - d_B(\psi(\gamma)(b), \pi_B(\overline{F})) - d_B(\psi(\gamma)(b'), \pi_B(\overline{F}')) \\ & \geq (d_B(b, b')/k - k - 2H') - 2H' \\ & = d_B(b, b')/k - k - 4H'. \end{aligned}$$

On the other hand, we also have

$$\begin{aligned} d_B(\psi(\gamma)(b), \psi(\gamma)(b')) & = d_B(\gamma((b, 0)), \gamma((b', 0))) \\ & \leq kd((b, 0), (b', 0)) + k \\ & \leq kd_B(b, b') + k. \end{aligned}$$

Having $(k + 2H')$ -dense image, the map $\psi(\gamma): B \rightarrow B$ is therefore a (k', k') -quasi-isometry with k' -dense image, where k' is a universal constant. It is now easy to

show that the map $\gamma \mapsto \psi(\gamma)$ defines a quasi-action of Γ on B . Moreover, up to increasing k' we may assume that such a quasi-action is k' -cobounded. From the way the action of Γ on B was defined, we also have that, for every $\gamma \in \Gamma$ and every component O of ∂B , there exists a component O' of ∂B such that the Hausdorff distance between $\psi(\gamma)(O)$ and O' is bounded by H . In order to simplify notations, we will as usual denote $\psi(\gamma)$ simply by γ .

Recall that $m = n - d$ is the dimension of the neutered space B , and let G be the isometry group of (B, d_B) . Every element of G is the restriction to B of an isometry of the whole hyperbolic space \mathbb{H}^m containing B . We will denote by $\text{Comm}(G)$ the *commensurator* of G in $\text{Isom}(\mathbb{H}^m)$, *i.e.* the group of those elements $h \in \text{Isom}(\mathbb{H}^m)$ such that the intersection $G \cap (hGh^{-1})$ has finite index both in G and in hGh^{-1} .

We are now in a position to use a deep result due to Schwartz (see [Sc, Lemma 6.1]), which in our context can be stated as follows:

THEOREM 10.2 ([Sc]). *There exists a universal constant β such that the following condition holds: for every $\gamma \in \Gamma$ a unique isometry $\theta(\gamma) \in \text{Isom}(\mathbb{H}^m)$ exists such that $d_{\mathbb{H}}(\gamma(x), \theta(\gamma)(x)) \leq \beta$ for every $x \in B$, where $d_{\mathbb{H}}$ denotes the hyperbolic distance on \mathbb{H}^m . Moreover, for every $\gamma \in \Gamma$ the isometry $\theta(\gamma)$ belongs to $\text{Comm}(G)$, and the resulting map $\theta: \Gamma \rightarrow \text{Comm}(G)$ is a group homomorphism.*

In the next few sections, we will analyze the kernel and image of the morphism θ , in order to extract information about the structure of Γ .

10.2. The image of θ

From now on we denote by $\Lambda < \text{Isom}(\mathbb{H}^m)$ the image of the homomorphism θ . Our next goal is to show that Λ is commensurable with $\pi_1(N)$. It is a result of Margulis that a non-uniform lattice in $\text{Isom}(\mathbb{H}^m)$ is arithmetic if and only if it has infinite index in its commensurator (see [Zi]). As a result, things would be quite a bit easier if N were assumed to be non-arithmetic. To deal with the general case, we will again use results (and techniques) from [Sc]. Note that, at this stage, we don't even know that Λ is a discrete subgroup of $\text{Isom}(\mathbb{H}^m)$.

From now on, unless otherwise stated, we will consider the Hausdorff distance of subsets of \mathbb{H}^m with respect to the hyperbolic metric $d_{\mathbb{H}}$ on \mathbb{H}^m . We denote by $P \subseteq \partial\mathbb{H}^m$ the set of all the basepoints of horospheres in ∂B . As an immediate corollary of Theorem 10.2 we get the following:

LEMMA 10.3. *For every $\alpha \in \Lambda$ and every horosphere $O \subseteq \partial B$ there exists a unique horosphere $O' \subseteq \partial B$ such that the Hausdorff distance between $\alpha(O)$ and O' is at most the universal constant $H + \beta$. In particular, the group Λ acts on P .*

LEMMA 10.4. *The action of Λ on the set P has a finite number of orbits, and every element of Λ which fixes a point in P is parabolic.*

PROOF. Fix a point $b \in B$. Let A be the set of boundary components of B whose hyperbolic distance from b is $\leq k'(H + k') + k'$. The set A is finite, and define P_0 to be the (finite) set of basepoints corresponding to the horospheres in the set A . We will prove that P_0 contains a set of representatives for the action of Λ on P .

So taking an arbitrary $p \in P$, let O be the corresponding component of ∂B , and fix a point $y \in O$. Since the quasi-action of Γ on B is k' -cobounded, there

exists $\gamma \in \Gamma$ such that $d_{\mathbb{H}}(\gamma(b), y) \leq d_B(\gamma(b), y) \leq k'$. We know that there exists a component O' of ∂B based at $p' \in P$ such that $\gamma(O')$ is at Hausdorff distance bounded by H from O . It follows that $\gamma(O')$ contains a point at distance at most H from y , and this in turn implies that O' belongs to A , so p' belongs to P_0 . Moreover, the horosphere $\theta(\gamma)(O')$ is at bounded Hausdorff distance from O , giving us $\theta(\gamma)(p') = p$. So p belongs to the Λ -orbit of a point in P_0 , completing the first part of the Lemma.

Now assume $p \in P$ is fixed by an element $\alpha \in \Lambda$, and let O be the connected component of ∂B corresponding to p . Since $\alpha(p) = p$, the horosphere $\alpha(O)$ is also based at the point p . It easily follows that the Hausdorff distance between O and $\alpha^n(O)$ equals n times the Hausdorff distance between O and $\alpha(O)$. Since $\alpha^n \in \Lambda$ for every $n \in \mathbb{N}$, if such a distance were positive, then for sufficiently large n the Hausdorff distance from O to $\alpha^n(O)$ would exceed the uniform constant $H + \beta$, contradicting Lemma 10.3. We conclude $\alpha(O) = O$, so α is parabolic. \square

Now let $P_0 = \{p_1, \dots, p_j\} \subseteq P$ as in Lemma 10.4 be a finite set of representatives for the action of Λ on P . For every $i = 1, \dots, j$ let O_i be the component of ∂B based at p_i , and let \widehat{O}_i be the horosphere contained in the horoball bounded by O_i and having Hausdorff distance $H + \beta$ from O_i . We let $\widehat{\mathcal{O}}$ be the set of horospheres obtained by translating $\widehat{O}_1, \dots, \widehat{O}_j$ by all the elements of Λ , and we denote by \widehat{B} the complement in \mathbb{H}^m of the union of the horoballs bounded by elements in $\widehat{\mathcal{O}}$. By construction the set \widehat{B} is Λ -invariant, and since all the stabilizers of points in P are parabolic, for every $p \in P$ there exists exactly one horosphere in $\widehat{\mathcal{O}}$ based at p . Let $R > 0$ be the minimal distance between distinct connected components of ∂B . Take $\widehat{O} \in \widehat{\mathcal{O}}$ and let O be the corresponding boundary component of B . By definition there exist $i \in \{1, \dots, j\}$ and an element $\alpha \in \Lambda$ such that $\widehat{O} = \alpha(\widehat{O}_i)$. Recall now that the Hausdorff distance between $\alpha(O_i)$ and O is bounded by $H + \beta$. Together with our choice for the construction of \widehat{O}_i , this implies that \widehat{O} is contained in the horoball bounded by O , and the Hausdorff distance between \widehat{O} and O is bounded by $2(H + \beta)$. As a consequence we easily deduce the following:

LEMMA 10.5. *The set \widehat{B} is Λ -invariant and is such that*

$$B \subseteq \widehat{B} \subseteq N_{2(H+\beta)}(B)$$

(where regular neighbourhoods are considered with respect to the hyperbolic metric $d_{\mathbb{H}}$). Moreover, if $\widehat{O}, \widehat{O}'$ are distinct elements of $\widehat{\mathcal{O}}$, then the Hausdorff distance between the horoballs bounded by \widehat{O} and \widehat{O}' is at least R (in particular, such horoballs are disjoint).

We are now ready to prove the following:

PROPOSITION 10.6. *The group Λ is a non-uniform lattice in $\text{Isom}(\mathbb{H}^m)$, and admits \widehat{B} as associated neutered space.*

PROOF. We begin by showing that Λ is discrete. Since N has finite volume, the set P is dense in $\partial \mathbb{H}^m$, so we may find horospheres $\widehat{O}_1, \dots, \widehat{O}_{m+1}$ in $\partial \widehat{B}$ with basepoints p_1, \dots, p_{m+1} such that $\{p_1, \dots, p_{m+1}\}$ is not contained in the trace at infinity of any hyperbolic hyperplane of \mathbb{H}^m . In particular, if $\alpha \in \text{Isom}(\mathbb{H}^m)$ is such that $\alpha(p_i) = p_i$ for every $i = 1, \dots, m+1$, then $\alpha = \text{Id}$.

Recall that the minimal distance between distinct connected components of $\partial\widehat{B}$ is bounded from below by the constant $R > 0$. Choose $x_i \in O_i$ for $i = 1, \dots, m+1$ and set

$$U = \{\alpha \in \text{Isom}(\mathbb{H}^m) \mid d_{\mathbb{H}}(\alpha(x_i), x_i) < R \text{ for every } i = 1, \dots, m+1\}.$$

Then U is an open neighbourhood of the identity in $\text{Isom}(\mathbb{H}^m)$; let us compute the intersection $\Lambda \cap U$. If $\alpha \in \Lambda$, we have that α permutes the component of $\partial\widehat{B}$. If we also assume $\alpha \in U$, then α moves each of the horospheres O_i at most R , which forces $\alpha(O_i) = O_i$, whence $\alpha(p_i) = p_i$, for each $i = 1, \dots, m+1$. As noted above, this implies $\alpha = \text{Id}$, and $\Lambda \cap U = \{\text{Id}\}$. But this implies Λ is a discrete subgroup.

Next we verify that Λ has finite co-volume. Since \widehat{B} is contained in the $2(H+\beta)$ -neighbourhood of B , there exists a Γ -orbit which is $(k' + 2H + 2\beta)$ -dense in \widehat{B} , and this immediately implies that there exists a Λ -orbit which is $(k' + 2H + 3\beta)$ -dense in \widehat{B} . It follows that the quotient orbifold \widehat{B}/Λ is compact. By Lemma 10.4, such an orbifold has a finite number V_1, \dots, V_j of boundary components. Let \widehat{O}_j be the boundary component of \widehat{B} projecting onto V_j . Since elements of Λ permute the boundary components of \widehat{B} , if $\alpha \in \Lambda$ is such that $\alpha(\widehat{O}_j) \cap \widehat{O}_j \neq \emptyset$, then $\alpha(\widehat{O}_j) = \widehat{O}_j$, so α belongs to the stabilizer Λ_j of the basepoint of \widehat{O}_j . Being a closed subset of the compact quotient \widehat{B}/Λ , the set $V_j = \widehat{O}_j/\Lambda_j$ is also compact. If $W_j \subseteq \mathbb{H}^m$ is the horoball bounded by \widehat{O}_j , it follows that the quotient W_j/Λ_j has finite volume. Since $(\bigcup_{i=1}^j W_j) \cup \widehat{B}$ projects surjectively onto \mathbb{H}^m/Λ , we conclude that \mathbb{H}^m/Λ has finite volume, and we have verified that Λ is a non-uniform lattice. \square

COROLLARY 10.7. *The group Λ is commensurable with $\pi_1(N)$.*

PROOF. Since $B \subseteq \widehat{B} \subseteq N_{2(H+\beta)}(B)$, the spaces B and \widehat{B} , when endowed with their path distances, are quasi-isometric. Since $\pi_1(N)$ acts properly and cocompactly on B and Λ acts properly and cocompactly on \widehat{B} , by Milnor-Svarc's Lemma this ensures that Λ is quasi-isometric to $\pi_1(N)$. The conclusion now follows from [Sc, Corollary 1.3], since both $\pi_1(N)$ and Λ are non-uniform lattices in $\text{Isom}(\mathbb{H}^m)$. \square

10.3. The kernel of θ

Having obtained an understanding of the image of θ , we now turn to studying the kernel.

LEMMA 10.8. *The group $\ker \theta$ is finitely generated and quasi-isometric to \mathbb{Z}^d . Moreover, it is quasi-isometrically embedded in Γ .*

PROOF. Let $F = \{b\} \times \mathbb{R}^d \subseteq \widetilde{M}$ be a fixed fiber of \widetilde{M} , set $x_0 = (b, 0) \in F$ and observe that there exists $\beta' > 0$ such that if $\gamma \in \ker \theta$ then $\gamma(x_0) \in N_{\beta'}(F)$ (we may take as β' the smallest number such that in the base B every $d_{\mathbb{H}}$ -ball of radius β is contained in a d_B -ball of radius β'). For $\gamma \in \ker \theta$, $x \in F$, we denote by $\alpha(\gamma, x) \in F$ a point such that $d(\alpha(\gamma, x), \gamma(x)) \leq \beta'$. It is not difficult to see that the resulting map $\alpha: \ker \theta \times F \rightarrow F$ defines a quasi-action. Since the fiber F is isometric to \mathbb{R}^d (and hence quasi-isometric to \mathbb{Z}^d), Lemma 1.4 tells us the first statement would follow provided we can show that α is cobounded, *i.e.* that the orbit of x_0 is Q -dense in F for some Q .

First observe that if $\gamma \in \Lambda$ is such that $\gamma(x_0) \in N_{\beta'}(F)$, then $\theta(\gamma)$ moves b a universally bounded distance from itself, so discreteness of Λ implies that $\theta(\gamma)$ belongs to a fixed finite subset $A \subseteq \Lambda$. For every $a \in A$ we choose an element $\gamma_a \in \Gamma$ such that $\theta(\gamma_a) = a$ and we set $M = \max\{d(x_0, \gamma_a^{-1}(x_0)), a \in A\}$. Now, for each point $p \in F$ there exists $\gamma \in \Gamma$ such that $d(\gamma(x_0), p) \leq k$. Then, if $\theta(\gamma) = a \in A$ we have that $\gamma\gamma_a^{-1} \in \ker \theta$ and

$$\begin{aligned} d((\gamma\gamma_a^{-1})(x_0), p) &\leq d(\gamma(\gamma_a^{-1}(x_0)), p) + k \\ &\leq d(\gamma(\gamma_a^{-1}(x_0)), \gamma(x_0)) + d(\gamma(x_0), p) + k \\ &\leq kd(\gamma_a^{-1}(x_0), x_0) + 3k \\ &\leq M + 3k \end{aligned}$$

so $d(\alpha(\gamma\gamma_a^{-1}, x_0), p) \leq M + 3k + \beta'$. We have thus proved that α is cobounded, and from Lemma 1.4 we can now deduce that $\ker \theta$ is finitely-generated and quasi-isometric to F (whence to \mathbb{Z}^d) via the map

$$j_{x_0}: \ker \theta \rightarrow F, \quad j_{x_0}(\gamma) = \alpha(\gamma, x_0).$$

Let us now prove that $\ker \theta$ is quasi-isometrically embedded in Γ . Let $\varphi: \Gamma \rightarrow \widetilde{M}$, $\psi: \widetilde{M} \rightarrow \Gamma$ be the quasi-isometries introduced in Section 1.3, and let $i: F \rightarrow \widetilde{M}$ be the inclusion. Also choose k'' large enough, so that ψ is a (k'', k'') -quasi-isometry and $d(\psi(\varphi(\gamma)), \gamma) \leq k''$ for every $\gamma \in \Gamma$. Since F is totally geodesic in \widetilde{M} , the inclusion i defines an isometric embedding of F into M , hence the composition of quasi-isometric embeddings $\psi \circ i \circ j_{x_0}: \ker \theta \rightarrow \Gamma$ is also a quasi-isometric embedding. In order to conclude, it is now sufficient to show that the inclusion of $\ker \theta$ into Γ stays at bounded distance from $\psi \circ i \circ j_{x_0}$.

Keeping the notation from Section 1.3 (and recalling that, in the proof above, we denoted by $\gamma(x_0)$ the point $\varphi(\gamma \cdot \psi(x_0))$), for every $\gamma \in \ker \theta$ we have the series of inequalities:

$$\begin{aligned} d(\psi(i(j_{x_0}(\gamma))), \gamma) &= d(\psi(\alpha(\gamma, x_0)), \gamma) \\ &\leq d(\psi(\alpha(\gamma, x_0)), \psi(\gamma(x_0))) + d(\psi(\gamma(x_0)), \gamma) \\ &\leq k''\beta' + k'' + d(\psi(\varphi(\gamma \cdot \psi(x_0))), \gamma) \\ &\leq k''\beta' + 2k'' + d(\gamma \cdot \psi(x_0), \gamma) \\ &= k''\beta' + 2k'' + d(\psi(x_0), 1_\Gamma), \end{aligned}$$

where the last equality is due to the Γ -invariance of any word metric on Γ , and this concludes the proof. \square

We now need the following fundamental result by Gromov:

THEOREM 10.9 ([Gr1]). *A finitely generated group quasi-isometric to \mathbb{Z}^d contains a finite index subgroup isomorphic to \mathbb{Z}^d .*

By Theorem 10.9, $\ker \theta$ contains a finite index subgroup K isomorphic to \mathbb{Z}^d . Being finitely generated, $\ker \theta$ contains only a finite number of subgroups having the same index as K . The intersection of all such subgroups has finite index in K and is characteristic in $\ker \theta$. Therefore, up to replacing K with one of its finite index subgroups, we can assume that K is characteristic in $\ker \theta$, hence normal in Γ . By construction, the quotient Γ/K is a finite extension of $\Lambda = \Gamma/\ker \theta$. By Corollary 10.7, there exists a finite index subgroup Λ' of Λ such that $\Lambda' \cong \pi_1(N')$

for some finite-sheeted covering N' of N . Let us set $\Gamma' = \theta^{-1}(\Lambda')$ and $\Delta = \Gamma'/K$. Then, we have the following exact sequences:

$$(10.1) \quad \begin{aligned} 1 &\longrightarrow \mathbb{Z}^d \xrightarrow{j} \Gamma' \xrightarrow{\theta} \Gamma'/K = \Delta \longrightarrow 1, \\ 1 &\longrightarrow F \longrightarrow \Delta \longrightarrow \pi_1(N') \longrightarrow 1, \end{aligned}$$

where $K = j(\mathbb{Z}^d)$, and F is finite.

10.4. Analyzing the exact sequence

In order to conclude the proof of Theorem 0.13, all that remains is to show that the sequence (10.1) is virtually central, *i.e.* that $K = j(\mathbb{Z}^d)$ is contained in the center of a finite-index subgroup of Γ' . To this aim, we will introduce the notion of *translation number*, and follow a strategy already described in [Gr2, GeSh] (see also [AlBr, KILe]).

Let G be a finitely-generated group with finite set of generators A , and for every $g \in G$ let us denote by $|g|_A$ the distance between g and the identity of G in the Cayley graph of G relative to A . The *translation number* of g is then given by the non-negative number

$$\tau_{G,A}(g) = \lim_{n \rightarrow \infty} \frac{|g^n|_A}{n}$$

(the fact that such a limit exists follows from the inequality $|g^{m+n}|_A \leq |g^m|_A + |g^n|_A$, which holds for every $g \in G$, $m, n \in \mathbb{N}$). In what follows we will need the following elementary properties of the translation number:

- (1) $\tau_{G,A}(ghg^{-1}) = \tau_{G,A}(h)$ for every $g, h \in G$;
- (2) if G is free abelian and A is a basis of G , then $\tau_{G,A}(g) = |g|_A$ for every $g \in G$;
- (3) let G be a subgroup of G' and A, A' be finite set of generators for G, G' ; if the inclusion $i: G \rightarrow G'$ is a (λ, ε) -quasi-isometric embedding (with respect to the metrics defined on G, G' by A, A'), then for every $g \in G$ we have

$$\lambda^{-1}\tau_{G,A}(g) \leq \tau_{G',A'}(g) \leq \lambda\tau_{G,A}(g).$$

Let us now consider the exact sequence (10.1). We fix an element $x \in \Delta$, and let $y \in \Gamma'$ be such that $\theta(y) = x$. Since $K = j(\mathbb{Z}^d)$ is normal in Γ' and abelian, the automorphism $\alpha(x): K \rightarrow K$ defined by $\alpha(x)(k) = y \cdot k \cdot y^{-1}$ is well-defined (in particular, it does not depend on the choice of y), so the map $\alpha: \Delta \rightarrow \text{Aut}(K)$ is a well-defined homomorphism of groups.

Now let $A' \subseteq \Gamma'$ be a finite set of generators and let $A = \{k_1, \dots, k_d\}$ be a free basis of K . For every $x \in \Delta$, $i = 1, \dots, d$, the element $\alpha(x)(k_i)$ is obviously conjugate to k_i in Γ' , so by property (1) above we have

$$(10.2) \quad \tau_{\Gamma',A'}(\alpha(x)(k_i)) = \tau_{\Gamma',A'}(k_i).$$

Since K is a finite-index subgroup of $\ker \theta$ and Γ' is a finite-index subgroup of Γ , by Lemma 10.8 the inclusion of K in Γ' is a quasi-isometric embedding. So by property (3) of the translation number there exists $\lambda > 0$ such that

$$(10.3) \quad \tau_{K,A}(\alpha(x)(k_i)) \leq \lambda\tau_{\Gamma',A'}(\alpha(x)(k_i)), \quad \tau_{\Gamma',A'}(k_i) \leq \lambda\tau_{K,A}(k_i) = \lambda.$$

Putting together property (2) of the translation number with equations (10.2) and (10.3) we finally obtain

$$|\alpha(x)(k_i)|_A = \tau_{K,A}(\alpha(x)(k_i)) \leq \lambda_{\Gamma',A'}(\alpha(x)(k_i)) = \lambda_{\Gamma',A'}(k_i) \leq \lambda^2$$

for every $x \in \Delta$, $i = 1, \dots, d$. This implies that the orbit of each k_i under the action of $\alpha(\Delta)$ is finite, so the homomorphism $\alpha: \Delta \rightarrow \text{Aut}(K)$ has finite image, and $\ker \alpha$ has finite index in $\pi_1(N')$. Moreover, from the very definition of α , it follows that K is contained in the center of $\theta^{-1}(\ker \alpha)$, which is a finite-index subgroup of Γ' , whence of Γ . We conclude that, up to replacing Γ' with another finite-index subgroup of Γ , we can suppose that the extension described by the sequence (10.1) is central. The proof of Theorem 0.13 is now complete.

REMARK 10.10. One may wonder whether the assumption that Γ (whence Γ') is quasi-isometric to $\pi_1(N') \times \mathbb{Z}^d$ implies that sequence (10.1) above virtually splits. The following construction shows that this is not true in general.

Let N be a hyperbolic 3-manifold with $k \geq 1$ cusps and second Betti number $b_2 > k$ (it is not difficult to construct such a manifold, for example by considering suitable link complements in the connected sum of several copies of $S^2 \times S^1$), and let \widehat{N} be obtained by Dehn filling all the cusps of N , in such a way that \widehat{N} is hyperbolic. An easy argument using a Mayer-Vietoris sequence shows that a 2-class $c_N \in H_2(N; \mathbb{Z})$ exists such that the element $i_*(c_N) \in H_2(\widehat{N}; \mathbb{Z})$ has infinite order, where $i: N \rightarrow \widehat{N}$ is the natural inclusion. Thanks to the Universal Coefficient Theorem, a coclass $\omega \in H^2(\widehat{N}; \mathbb{Z})$ exists such that $\omega(i_*(c_N)) = 1$ (here and henceforth we denote by $\omega(i_*(c_N))$ the number $\langle \omega, i_*(c_N) \rangle$, where $\langle \cdot, \cdot \rangle: H^2(\widehat{N}; \mathbb{Z}) \times H_2(\widehat{N}; \mathbb{Z}) \rightarrow \mathbb{Z}$ is the Kronecker pairing).

Since N and \widehat{N} have contractible universal coverings, we have natural isomorphisms $H_2(N; \mathbb{Z}) \cong H_2(\pi_1(N); \mathbb{Z})$, $H_2(\widehat{N}; \mathbb{Z}) \cong H_2(\pi_1(\widehat{N}); \mathbb{Z})$, $H^2(N; \mathbb{Z}) \cong H^2(\pi_1(N); \mathbb{Z})$, $H^2(\widehat{N}; \mathbb{Z}) \cong H^2(\pi_1(\widehat{N}); \mathbb{Z})$. Abusing notation, we will also denote by $c_N \in H_2(\pi_1(N); \mathbb{Z})$, $i_*(c_N) \in H_2(\pi_1(\widehat{N}); \mathbb{Z})$, $\omega \in H^2(\pi_1(\widehat{N}); \mathbb{Z})$ the elements corresponding to the (co)classes introduced above. The inclusion $i: N \hookrightarrow \widehat{N}$ induces a morphism $i^*: H^2(\pi_1(\widehat{N}); \mathbb{Z}) \rightarrow H^2(\pi_1(N); \mathbb{Z})$.

We now consider the central extension

$$1 \rightarrow \mathbb{Z} \rightarrow G \rightarrow \pi_1(N) \rightarrow 1$$

associated to the coclass $i^*(\omega) \in H^2(\pi_1(N); \mathbb{Z})$. On one hand, since $\pi_1(\widehat{N})$ is Gromov-hyperbolic, by [NeRe] the coclass $\omega \in H^2(\pi_1(\widehat{N}); \mathbb{Z})$ admits a bounded representative, so $i^*(\omega)$ is also bounded. By [NeRe], this implies that G is quasi-isometric to $\pi_1(N) \times \mathbb{Z}$. On the other hand, we have $i^*(\omega)(c_N) = \omega(i_*(c_N)) = 1$, so $i^*(\omega)$ has infinite order in $H^2(N; \mathbb{Z})$, and this proves that the above sequence does not virtually split.

10.5. Pieces with quasi-isometric fundamental groups

The following proposition provides a necessary and sufficient condition for two pieces of graph manifolds to have quasi-isometric fundamental groups.

PROPOSITION 10.11. *Let $n \geq 3$ be fixed, and, for $i = 1, 2$, let N_i be a complete finite-volume hyperbolic n_i -manifold with toric cusps, $n_i \geq 3$. If $\pi_1(N_1 \times T^{n-n_1}) = \pi_1(N_1) \times \mathbb{Z}^{n-n_1}$ is quasi-isometric to $\pi_1(N_2 \times T^{n-n_2}) = \pi_1(N_2) \times \mathbb{Z}^{n-n_2}$, then $n_1 = n_2$ and N_1 is commensurable with N_2 .*

PROOF. Let us set $G_i = \pi_1(N_i) \times \mathbb{Z}^{n-n_i}$. By Theorem 0.13, since G_1 is quasi-isometric to $\pi_1(N_2) \times \mathbb{Z}^{n-n_2}$, there exist a finite index subgroup G'_1 of G_1 , a group Δ and a finite group F which fit in the following short exact sequences:

$$\begin{aligned} 1 &\longrightarrow \mathbb{Z}^{n-n_2} \xrightarrow{j} G'_1 \longrightarrow \Delta \longrightarrow 1, \\ 1 &\longrightarrow F \longrightarrow \Delta \longrightarrow \pi_1(N'_2) \longrightarrow 1, \end{aligned}$$

where N'_2 is a finite-sheeted covering of N_2 . Moreover, $j(\mathbb{Z}^{n-n_2})$ lies in the center of G'_1 .

Let $Z(G_1)$ (resp. $Z(G'_1)$) be the center of G_1 (resp. of G'_1). We claim that $Z(G'_1) = Z(G_1) \cap G'_1$. The inclusion \supseteq is obvious. Moreover, if $p_1: G_1 \rightarrow \pi_1(N_1)$ is the projection on the first factor, then $p_1(G'_1)$ is a finite-index subgroup of $\pi_1(N_1)$. Since any finite-index subgroup of $\pi_1(N_1)$ has trivial center, this implies that any element $(\gamma, w) \in G'_1 \subseteq G_1 = \pi_1(N_1) \times \mathbb{Z}^{n-n_1}$ which commutes with all the elements of G'_1 must satisfy $\gamma = 1$ in $\pi_1(N_1)$. We conclude that $(\gamma, w) \in Z(G_1)$, as claimed.

This implies that $j(\mathbb{Z}^{n-n_2}) \subseteq Z(G'_1) \subseteq Z(G_1) \cong \mathbb{Z}^{n-n_1}$, so $n_1 \leq n_2$ by injectivity of j . Interchanging the roles of G_1 and G_2 we also get $n_2 \leq n_1$, forcing $n_1 = n_2$.

Since $Z(G'_1) = Z(G_1) \cap G'_1$, the quotient $G'_1/Z(G'_1)$ is isomorphic to a finite-index subgroup of $G_1/Z(G_1)$, which is in turn isomorphic to $\pi_1(N_1)$. In particular, $G'_1/Z(G'_1)$ is quasi-isometric to $\pi_1(N_1)$. Moreover, since $n_1 = n_2$ the groups $j(\mathbb{Z}^{n-n_2})$ and $Z(G'_1)$ share the same rank, and this implies that $j(\mathbb{Z}^{n-n_2})$ is a finite-index subgroup of $Z(G'_1)$, so that $\Delta \cong G'_1/j(\mathbb{Z}^{n-n_2})$ is quasi-isometric to $G'_1/Z(G'_1)$, whence to $\pi_1(N_1)$. On the other hand, since Δ is a finite extension of $\pi_1(N'_2)$ and $\pi_1(N'_2)$ is of finite index in $\pi_1(N_2)$, the group Δ is quasi-isometric to $\pi_1(N_2)$ too, so $\pi_1(N_1)$ and $\pi_1(N_2)$ are quasi-isometric to each other. The conclusion now follows from [Sc]. \square

Quasi isometry rigidity, II

The first aim of this section is the proof of Theorem 0.14, which we recall here:

THEOREM. *Let M be an irreducible graph n -manifold obtained by gluing the pieces $V_i = \bar{N}_i \times T^{d_i}$, $i = 1, \dots, k$. Let Γ be a group quasi-isometric to $\pi_1(M)$. Then either Γ itself, or a subgroup of Γ of index two, is isomorphic to the fundamental group of a graph of groups satisfying the following conditions:*

- every edge group contains \mathbb{Z}^{n-1} as a subgroup of finite index;
- for every vertex group Γ_v there exist $i \in \{1, \dots, k\}$, a finite-sheeted covering N' of N_i and a finite-index subgroup Γ'_v of Γ_v that fits into the exact sequences

$$1 \longrightarrow \mathbb{Z}^{d_i} \xrightarrow{j} \Gamma'_v \longrightarrow \Delta \longrightarrow 1,$$

$$1 \longrightarrow F \longrightarrow \Delta \longrightarrow \pi_1(N') \longrightarrow 1,$$

where F is a finite group, and $j(\mathbb{Z}^{d_i})$ is contained in the center of Γ'_v .

Throughout this chapter we denote by M an irreducible graph manifold with universal covering \widetilde{M} , and by Γ a finitely generated group quasi-isometric to $\pi_1(M)$. As discussed in Section 1.3, a quasi-isometry between Γ and $\pi_1(M)$ induces a k -cobounded k -quasi-action h of Γ on \widetilde{M} for some $k \geq 1$, which will from now on be fixed. Henceforth, for every $\gamma \in \Gamma$ we will denote simply by γ the quasi-isometry $h(\gamma): \widetilde{M} \rightarrow \widetilde{M}$.

11.1. From quasi-actions to actions on trees

Let (\widetilde{M}, p, T) be the triple which endows \widetilde{M} with the structure of a tree of spaces (see Section 2.3). Building on the results proved in Chapter 9, we wish to define an action of Γ on T . Fix $\gamma \in \Gamma$. By Propositions 9.16 and 9.17, if v_1, e_1 are a vertex and an edge corresponding respectively to a chamber C_1 and a wall W_1 , then there exist a unique chamber C_2 at finite Hausdorff distance from $\gamma(C_1)$ and a unique wall W_2 at finite Hausdorff distance from $\gamma(W_1)$. We will denote by $\gamma(v_1), \gamma(e_1)$ the vertex and the edge corresponding respectively to C_2 and W_2 . Again, by Proposition 9.17, if W_1 is adjacent to C_1 then W_2 is adjacent to C_2 , which gives us the following:

PROPOSITION 11.1. *The map $\gamma: T \rightarrow T$ just defined provides a simplicial automorphism of T .*

In what follows, when saying that a group G acts on a tree T' we will always mean that G acts on T' by simplicial automorphisms. Recall that G acts on T' without inversions if no element of G switches the endpoints of an edge of T' . We wish to apply the following fundamental result from Bass-Serre theory (see [Se]):

THEOREM 11.2. *Suppose G acts on a tree T' without inversions. Then G is isomorphic to the fundamental group of a graph of groups supported by the graph \mathcal{G} with set of vertices V and set of edges E . If G_v , $v \in V$, and G_e , $e \in E$, are the vertex and edge groups of the graph of groups, then:*

- (1) \mathcal{G} is the quotient of T' by the action of G .
- (2) For each $v \in V$, the group G_v is isomorphic to the stabilizer of a vertex of T' projecting to v .
- (3) For each $e \in E$, the group G_e is isomorphic to the stabilizer of an edge of T' projecting to e .

Now the action of Γ on T described in Proposition 11.1 might include some inversions. However, Bass has shown in [Ba, Proposition 6.3] that the group $\text{Aut}(T)$ of all simplicial automorphisms of T contains a subgroup $\text{Aut}^0(T)$, of index at most two, which consists solely of elements that act without inversions. This immediately implies:

LEMMA 11.3. *If G acts on the tree T' , then either G acts without inversions or the action of G restricts to an action without inversions of a subgroup G^0 of index 2 in G .*

Applying Lemma 11.3 to the action of Γ on T , we conclude that a subgroup of index at most two acts on T without inversions. For the rest of this Chapter, we will denote by Γ^0 the subgroup of Γ having this property.

11.2. The action of Γ^0 on T

Recall that Γ^0 quasi-acts via (k, k) -quasi-isometries with k -dense image on \widetilde{M} , and, up to increasing the constant k , we may also assume that every Γ^0 -orbit is k -dense in \widetilde{M} . We denote by E the set of edges of T , and we suppose that for every wall W (resp. chamber C) and every $\gamma \in \Gamma^0$ the set $f(W)$ (resp. $f(C)$) has Hausdorff distance bounded by H from a wall (resp. a chamber) (see Propositions 9.16 and 9.17). We first show that the quotient of T by the action of Γ^0 is a finite graph.

LEMMA 11.4. *The action of Γ^0 on E has a finite number of orbits.*

PROOF. Fix a point $p \in \widetilde{M}$. The set A of those walls whose distance from p is less than $k(H + k) + 3k$ is finite. Let W be any wall, and fix a point $w \in W$. There exists $\gamma \in \Gamma^0$ such that $d(\gamma(p), w) \leq k$. We know that there exists a wall W' such that $\gamma(W')$ is at Hausdorff distance bounded by H from W . This implies that $\gamma(W')$ contains a point $\gamma(w')$, $w' \in W'$, at distance less than H from w . We can use this to estimate:

$$\begin{aligned} d(w', p) &\leq d(\gamma^{-1}(\gamma(w')), \gamma^{-1}(\gamma(p))) + 2k \\ &\leq kd(\gamma(w'), \gamma(p)) + 3k \\ &\leq k(H + k) + 3k, \end{aligned}$$

so $W' \in A$. As a result, the finite set of edges corresponding to walls in A contains a set of representatives for the action of Γ^0 on E . \square

11.3. Stabilizers of edges and vertices

If e (resp. v) is an edge (resp. a vertex) of T , then we denote by Γ_e^0 (resp. Γ_v^0) the stabilizer of e (resp. of v) in Γ^0 .

LEMMA 11.5. *For every edge e of T , the stabilizer Γ_e^0 is quasi-isometric to a wall. The stabilizer Γ_v^0 of a vertex v is quasi-isometric to the chamber corresponding to v .*

PROOF. Let us focus on proving the first statement, as the second statement follows from a very similar argument. Let $N_H(W)$ be the H -neighbourhood of the wall W corresponding to the edge $e \subseteq T$, and let $\varphi_e: \Gamma_e^0 \rightarrow N_H(W)$ be defined by $\varphi_e(\gamma) = \gamma(w)$, where $w \in W$ is a fixed basepoint. Let us first prove that $\varphi_e(\Gamma_e^0)$ is p -dense in $N_H(W)$ for some p . For each wall W_i , $i = 1, \dots, m$, in the orbit of W and having distance less than $k^2 + 2k + H$ from w , we choose $\gamma_i \in \Gamma^0$ such that $\gamma_i(W)$ has Hausdorff distance from W_i bounded by H . Let L be large enough so that $d(w, \gamma_i(w)) \leq L$ for every $i = 1, \dots, m$. Now pick any point $w' \in N_H(W)$. We know that there is $\gamma \in \Gamma^0$ (but not necessarily in Γ_e^0) such that $d(\gamma(w), w') \leq k$. It is not difficult to show that $\gamma^{-1}(W)$ has finite Hausdorff distance from one of the W_i 's, so there exists j such that $\gamma(W_j)$ is at finite Hausdorff distance from W . Then $\gamma \cdot \gamma_j \in \Gamma_e^0$, and we have the estimate:

$$\begin{aligned} d((\gamma\gamma_j)(w), w') &\leq d((\gamma\gamma_j)(w), \gamma(w)) + d(\gamma(w), w') \\ &\leq (d(\gamma(\gamma_j(w)), \gamma(w)) + k) + k \\ &\leq (kd(\gamma_j(w), w) + 2k) + k \\ &\leq kL + 3k. \end{aligned}$$

This implies that $\varphi_e(\Gamma_e^0)$ is $(kL + 3k)$ -dense in $N_H(W)$.

In order to apply Lemma 1.4 we now need to construct a quasi-action of Γ_e^0 on (W, d_W) , where d_W is the path-distance of W . With this goal in mind, for every $\gamma \in \Gamma_e^0$ and $x \in W$, we let $h_e(\gamma)(x)$ be a point in W such that $d(\gamma(x), h_e(\gamma)(x)) \leq H$. It is easily checked that the map $\gamma \mapsto h_e(\gamma)$ indeed defines a quasi-action of Γ_e^0 on (W, d) , where d is the restriction to W of the distance on \widetilde{M} . Moreover, the orbit of w under this quasi-action is $(kL + 3k + 2H)$ -dense in (W, d) . But since M is irreducible the identity map on W provides a quasi-isometry between (W, d) and the path metric space (W, d_W) , so h provides a quasi-action of Γ_e^0 on (W, d_W) , and the orbit of w is p -dense in (W, d_W) for some p . By Lemma 1.4, this implies that Γ_e^0 is finitely generated and quasi-isometric to (W, d_W) . \square

REMARK 11.6. Arguing as in the proof of Lemma 10.8, it is possible to prove that the stabilizers Γ_e^0, Γ_v^0 are quasi-isometrically embedded in Γ^0 .

Putting together Lemma 11.5 and Gromov's Theorem 10.9 we immediately get the following:

PROPOSITION 11.7. *If Γ_e^0 is the stabilizer of an edge $e \subseteq T$, then Γ_e^0 contains \mathbb{Z}^{n-1} as a subgroup of finite index.*

Theorem 0.14 is now a direct consequence of Theorem 11.2, Proposition 11.7 and Theorem 0.13.

11.4. Graph manifolds with quasi-isometric fundamental groups

We are now interested in analyzing when irreducible graph manifolds have quasi-isometric fundamental groups.

For $i = 1, 2$, let M_i be an irreducible graph manifold, and let us denote by T_i the tree corresponding to the decomposition of \widetilde{M}_i into chambers. We can label each vertex v of T_i as follows: if v corresponds to a chamber projecting in M onto a piece of the form $N \times T^d$, where N is a cusped hyperbolic manifold, then we label v with the commensurability class of N . The following result gives a necessary condition for M_1, M_2 to have quasi-isometric fundamental groups:

THEOREM 11.8. *Suppose the fundamental groups of M_1 and M_2 are quasi-isometric. Then T_1 and T_2 are isomorphic as labelled trees.*

PROOF. By Milnor-Svarc's Lemma, a quasi-isometry between $\pi_1(M_1)$ and $\pi_1(M_2)$ induces a quasi-isometry, say ψ , between the universal coverings \widetilde{M}_1 and \widetilde{M}_2 . By Proposition 9.17 (see also Subsection 1.3), such a quasi-isometry induces a simplicial isomorphism f_ψ between T_1 and T_2 . We will now show that such isomorphism preserve labels, thus proving the theorem.

Let v_1 be a vertex of T_1 corresponding to the chamber C_1 , and suppose that C_1 is the universal covering of $N_1 \times T^{d_1}$, where N_1 is a cusped hyperbolic manifold. Let C_2 be the chamber of \widetilde{M}_2 staying at finite Hausdorff distance from $\psi(C_1)$, let v_2 be the vertex of T_2 corresponding to C_2 , and suppose that C_2 projects into M_2 onto a piece of the form $N_2 \times T^{d_2}$, where N_2 is a cusped hyperbolic manifold. By construction, f_ψ takes v_1 onto v_2 , so we only need to check that the labels of v_1 and v_2 are equal, *i.e.* that N_1 is commensurable with N_2 . However, since M_1, M_2 are irreducible, the chamber C_i is quasi-isometrically embedded in M_i , and this implies that $\psi|_{C_1}$ stays at bounded distance from a quasi-isometry between C_1 and C_2 . By Milnor-Svarc's Lemma, it follows that $\pi_1(N_1) \times \mathbb{Z}^{d_1}$ is quasi-isometric to $\pi_1(N_2) \times \mathbb{Z}^{d_2}$, so N_1 is commensurable with N_2 by Proposition 10.11. \square

Observe that, in each dimension, there exist infinitely many commensurability classes of complete finite-volume hyperbolic manifolds with toric cusps (see [MRS]). Along with Theorem 11.8, this immediately allows us to deduce:

COROLLARY 11.9. *Suppose $n \geq 3$. Then, there exist infinitely many quasi-isometry classes of fundamental groups of irreducible graph n -manifolds.*

REMARK 11.10. Let us fix the notation as in Theorem 11.8. The following construction shows that the fact that T_1 and T_2 are isomorphic as labelled trees is **not** sufficient for ensuring that $\pi_1(M_1)$ and $\pi_1(M_2)$ are quasi-isometric.

Let N be a cusped hyperbolic 3-manifold with two toric cusps, let $\partial_1 \overline{N}$, $\partial_2 \overline{N}$ be the boundary components of the truncated manifold \overline{N} , and assume that the Euclidean structures induced by N on $\partial_1 \overline{N}$, $\partial_2 \overline{N}$ are not commensurable with each other. The fact that such a manifold exists is proved in [GHH] (we may take for example the manifold 7c 3548 in the census available at the address [www]). Furthermore, let N', N'' be non-commensurable 1-cusped hyperbolic 3-manifolds (for example, suitable hyperbolic knot complements), and consider the (obviously irreducible) graph manifolds M_1, M_2 defined as follows: M_1 is obtained by gluing \overline{N} with \overline{N}' along $\partial_1 \overline{N}$, and with \overline{N}'' along $\partial_2 \overline{N}$; M_2 is obtained by gluing \overline{N} with \overline{N}'' along $\partial_2 \overline{N}$, and with \overline{N}' along $\partial_1 \overline{N}$. Of course, the labelled trees associated to M_1 and M_2 are isomorphic.

On the other hand, a hypothetical quasi-isometry between $\pi_1(M_1)$ and $\pi_1(M_2)$ should induce a quasi-isometry of $\pi_1(N)$ into itself taking the cusp subgroup $\pi_1(\partial_1 \overline{N})$ to a set at finite Hausdorff distance from $\pi_1(\partial_2 \overline{N})$. By [Sc], this would imply that

$\pi_1(\partial_1 \overline{N})$ and $\pi_1(\partial_2 \overline{N})$ admit finite index subgroups that are conjugated by an isometry of \mathbb{H}^3 . As a consequence, the Euclidean structures induced by N on $\partial_1 \overline{N}$, $\partial_2 \overline{N}$ should be commensurable with each other, which would contradict our choices.

REMARK 11.11. In [BeNe], Behrstock and Neumann proved that the fundamental groups of any two closed 3-dimensional irreducible graph manifolds are quasi-isometric. This result could seem in contrast with the phenomenon exhibited by the previous construction. However, hyperbolic bases, in dimensions ≥ 3 , are much more rigid than hyperbolic surfaces with boundary. As a consequence, in higher dimensions there is no obvious counterpart for all the “stretchings” performed on thickened graphs in [BeNe].

Part 3

Concluding remarks

Examples not supporting locally CAT(0) metrics

We already saw a method in Section 2.5 for constructing graph manifolds which do not support any locally CAT(0) metric. The idea was to take a finite volume hyperbolic manifold N with at least two toric cusps, and glue together two copies of $N \times T^2$ in such a way that the fundamental group of the resulting graph manifold contains a non quasi-isometrically embedded abelian subgroup (see Proposition 2.19). This method could be used to produce infinitely many such examples in all dimensions ≥ 5 .

In this Chapter we provide some additional methods for constructing graph manifolds which do *not* support any locally CAT(0) metric. In Section 12.1 we show that certain S^1 -fiber-bundles over the double of cusped hyperbolic manifolds do *not* support locally CAT(0) metrics. This allows us to construct infinitely many new examples in each dimension ≥ 4 .

Section 12.2 is devoted to the construction of irreducible examples. We can produce infinitely many such examples in each dimension ≥ 4 .

For ease of notation, we will omit the coefficient ring in our cohomology groups, with the understanding that all homology and cohomology in this chapter is taken with coefficients in \mathbb{Z} .

12.1. Fiber bundles

In this section, we describe a construction providing graph manifolds which do not support any locally CAT(0) metrics. We start by recalling that principal S^1 -bundles over a manifold K are classified (topologically) by their Euler class in $H^2(K)$. The Euler class is the “primary obstruction” to the existence of a section, and satisfies the following two key properties:

Fact 1: The Euler class of a principal S^1 -bundle $S^1 \rightarrow K' \rightarrow K$ is zero if and only if $K' \cong K \times S^1$ (i.e. K' is the trivial S^1 -bundle).

Fact 2: If $f : L \rightarrow K$ is continuous, and $S^1 \rightarrow K' \rightarrow K$ is a principal S^1 -bundle, let $S^1 \rightarrow L' \rightarrow L$ be the pullback principal S^1 -bundle. Then $e(L') = f^*(e(K'))$, where $e(L'), e(K')$ denote the Euler classes of the respective S^1 -bundles, and $f^* : H^2(K; \mathbb{Z}) \rightarrow H^2(L; \mathbb{Z})$ is the induced map on the second cohomology.

Since the manifolds we will be considering arise as principal S^1 -bundles, we now identify a cohomological obstruction for certain principal S^1 -bundles to support a locally CAT(0) metric.

LEMMA 12.1. *Let K be a compact topological manifold supporting a locally CAT(0) metric, and let $S^1 \rightarrow K' \rightarrow K$ be a principal S^1 -bundle over K (so that K' is also compact). If K' supports a locally CAT(0) metric, then $e(K')$ has finite order in $H^2(K)$.*

PROOF. Since all spaces in the fibration are aspherical, the associated long exact sequence in homotopy degenerates to a single short exact sequence:

$$0 \rightarrow \mathbb{Z} \rightarrow \pi_1(K') \rightarrow \pi_1(K) \rightarrow 0.$$

As K' is compact, the action of $\pi_1(K')$ on the CAT(0) universal cover \tilde{K}' is by semi-simple isometries (i.e. for every $g \in \pi_1(K')$, there exists a $x \in \tilde{K}'$ satisfying $d(x, gx) \leq d(y, gy)$ for all $y \in \tilde{K}'$). Furthermore, $\pi_1(K')$ contains \mathbb{Z} as a normal subgroup. A well-known consequence of the Flat Torus theorem (see the discussion in [BrHa, pgs. 244-245]) implies that there exists a finite index subgroup $\Lambda \leq \pi_1(K')$ that centralizes the \mathbb{Z} -subgroup, i.e. we have:

$$\begin{array}{ccccc} \mathbb{Z} & \longrightarrow & \pi_1(K') & \longrightarrow & \pi_1(K) \\ \uparrow & & \uparrow & \text{Finite Index} & \\ \mathbb{Z} & \longrightarrow & \Lambda & \longrightarrow & \Lambda/\mathbb{Z} \end{array}$$

It is easy to see (by chasing the diagram) that there is an induced inclusion $\Lambda/\mathbb{Z} \hookrightarrow \pi_1(K)$ which is also of finite index. Let $L \rightarrow K$ be the finite cover corresponding to $\Lambda/\mathbb{Z} \hookrightarrow \pi_1(K)$, and $L' \rightarrow K'$ the cover corresponding to $\Lambda \hookrightarrow \pi_1(K')$. We now obtain the commutative diagram of principal bundles (see [BrHa, Theorem II.7.1-(5)]):

$$\begin{array}{ccccc} S^1 & \longrightarrow & K' & \longrightarrow & K \\ \uparrow & & \uparrow & & \uparrow \text{Finite Cover} \\ S^1 & \longrightarrow & L' & \longrightarrow & L \end{array}$$

where both the “top row” and the “bottom row” are principal S^1 -bundles. Now observe that the bottom row splits as a product, i.e. $L' \cong L \times S^1$. Indeed, this follows from the fact that Λ centralizes the \mathbb{Z} -factor, and splits as $\mathbb{Z} \oplus \Lambda/\mathbb{Z}$, while acting on the CAT(0) space \tilde{K}' . From **Fact 1**, this implies that $e(L') = 0 \in H^2(L)$. From **Fact 2**, and commutativity of the diagram, we get that $p^*(e(K')) = e(L') = 0$, where $p^* : H^2(K) \rightarrow H^2(L)$ is the map induced by the covering projection $p : L \rightarrow K$.

On the other hand, recall that there is a transfer map on cohomology $T : H^*(L) \rightarrow H^*(K)$ associated with any finite covering $p : L \rightarrow K$. This map has the property that $T \circ p^* : H^*(K) \rightarrow H^*(K)$ is just multiplication by the degree of the covering map. Hence if d denotes the degree of the covering map, we have that:

$$d \cdot e(K') = (T \circ p^*)(e(K')) = T(0) = 0 \in H^2(K)$$

implying that $e(K') \in H^2(K)$ is a torsion element, and completing the proof of the Lemma. \square

Keeping the notation from Section 2.1, let N be a finite volume, non-compact, hyperbolic manifold, with all cusps diffeomorphic to a torus times $[0, \infty)$, and let \bar{N} be the compact manifold obtained by “truncating the cusps”. Note that the boundary $\partial\bar{N}$ consists of a finite number of codimension one tori, and the inclusion $i : \partial\bar{N} \hookrightarrow \bar{N}$ induces the map $i^* : H^1(\bar{N}) \rightarrow H^1(\partial\bar{N})$ on the first cohomology. We will consider principal S^1 -bundles over the double $D\bar{N}$.

PROPOSITION 12.2. *Assume there exists a non-trivial cohomology class $\alpha \in H^1(\partial\bar{N})$ having the property that $\langle \alpha \rangle \cap i^*(H^1(\bar{N})) = 0 \subset H^1(\partial\bar{N})$. Then there exists a manifold M , which is topologically a principal S^1 -bundle over $D\bar{N}$, having the properties:*

- (1) M does **not** support any locally CAT(0) metric.
- (2) M is a graph manifold.

PROOF. It is well-known that the double $D\bar{N}$ supports a Riemannian metric of non-positive sectional curvature (see for example [ArFa, Theorem 1]). In view of Lemma 12.1, any principal S^1 -bundle whose Euler class has infinite order will **not** support any locally CAT(0) metric. Since every class in $H^2(D\bar{N})$ is realized as the Euler class of some principal S^1 -bundle, we just need to find a cohomology class of infinite order.

Consider the Mayer-Vietoris sequence in cohomology for the decomposition $D\bar{N} = \bar{N}_1 \cup_{\partial\bar{N}} \bar{N}_2$, where the \bar{N}_i are the two copies of \bar{N} . We have:

$$H^1(\bar{N}_1) \oplus H^1(\bar{N}_2) \xrightarrow{i} H^1(\partial\bar{N}) \xrightarrow{j} H^2(D\bar{N}) \longrightarrow H^2(\bar{N}_1) \oplus H^2(\bar{N}_2)$$

Now by hypothesis there exists an element $\alpha \in H^1(\partial\bar{N})$ having the property that $\langle \alpha \rangle \cap i^*(H^1(\bar{N})) = 0$. If i_1, i_2 denotes the inclusions of $\partial\bar{N}$ into \bar{N}_1, \bar{N}_2 , we have that the first map in the Mayer-Vietoris sequence above is given by $i := i_1^* - i_2^*$, and hence the non-trivial element $\alpha \in H^1(\partial\bar{N})$ has the property that $\langle \alpha \rangle \cap i(H^1(\bar{N}_1) \oplus H^1(\bar{N}_2)) = \{0\}$. In particular, since $H^1(\partial\bar{N})$ is torsion-free, the subgroup $j(H^1(\partial\bar{N})) \leq H^2(D\bar{N})$ contains an element of infinite order, namely $j(\alpha)$. Let M be the associated principal S^1 -bundle over $D\bar{N}$; from the discussion above, M cannot support any locally CAT(0) metric.

So to conclude, we just need to argue that M is a graph manifold. To see this, observe that M naturally decomposes as a union $M = M_1 \cup M_2$, where each M_i is the preimage of the respective \bar{N}_i under the canonical map $S^1 \rightarrow M \rightarrow D\bar{N} = \bar{N}_1 \cup_{\partial\bar{N}} \bar{N}_2$. We now show that the M_i are the pieces for the decomposition of M as a graph manifold. To do this, we need to understand the topology of the M_i .

From **Fact 2**, we can compute the Euler class of the bundles $S^1 \rightarrow M_i \rightarrow \bar{N}_i$ by looking at the image of $\alpha \in H^2(D\bar{N})$ under the maps $H^2(D\bar{N}) \rightarrow H^2(\bar{N}_i)$ induced by the inclusions $\bar{N}_i \hookrightarrow D\bar{N}$. But observe that these maps are exactly the ones appearing in the Mayer-Vietoris sequence:

$$H^1(\partial\bar{N}) \rightarrow H^2(D\bar{N}) \rightarrow H^2(\bar{N}_1) \oplus H^2(\bar{N}_2)$$

By exactness of the sequence, we immediately obtain that $\rho(j(\alpha)) = 0 \in H^2(\bar{N}_1) \oplus H^2(\bar{N}_2)$, and so the Euler class of both M_i is zero in the corresponding $H^2(\bar{N}_i)$. Applying **Fact 1**, we conclude that each M_i is the trivial S^1 -bundle over \bar{N}_i , *i.e.* each M_i is homeomorphic to $\bar{N}_i \times S^1$. Let us now endow each M_i with the smooth structure induced by the product $\bar{N}_i \times S^1$ of smooth manifolds. Now the only possible obstruction to M being a graph manifold lies in the gluing map between M_1 and M_2 being affine. However, if the gluing map is not affine, we can replace it by a homotopic affine diffeomorphism without affecting the Euler class of the corresponding principal S^1 -bundle (actually, if $n > 5$, we can replace the given gluing map by a C^0 -isotopic affine diffeomorphism without changing the topological type of the manifold M – see the discussion in Remark 2.4). Then M is indeed a graph manifold, and this concludes the proof of the Proposition. \square

In order to obtain the desired examples, we need to produce finite volume hyperbolic manifolds N so that the associated truncated \bar{N} satisfies:

- (1) all the boundary components of \bar{N} are diffeomorphic to tori, and
- (2) there exists a non-trivial element $\alpha \in H^1(\partial\bar{N})$ which satisfies

$$\langle \alpha \rangle \cap i^*(H^1(\bar{N})) = \{0\} \subset H^1(\partial\bar{N}).$$

The next step towards achieving this is to turn the cohomological condition (2) to a homological condition, as explained in the following Lemma.

LEMMA 12.3. *Let N be a finite volume hyperbolic manifold, so that the associated \bar{N} satisfies condition (1) above. Then N also satisfies condition (2) above if and only if $i_*: H_1(\partial\bar{N}) \rightarrow H_1(\bar{N})$ is not injective.*

PROOF. Since $H^1(\partial\bar{N})$ is a finitely generated torsion-free abelian group, property (2) above is equivalent to the fact that the index of $i^*(H^1(\bar{N}))$ in $H^1(\partial\bar{N})$ is infinite, so we need to prove that this last condition is in turn equivalent to the fact that $\ker i_* \neq \{0\}$.

For a torus T^k , the Kronecker pairing induces an isomorphism between $H^1(T^k)$ and $\text{Hom}(H_1(T^k), \mathbb{Z})$. Property (1) ensures that this duality extends to an isomorphism between $H^1(\partial\bar{N})$ and $\text{Hom}(H_1(\partial\bar{N}), \mathbb{Z})$. Moreover, it is easily seen that a subgroup H of $\text{Hom}(H_1(\partial\bar{N}), \mathbb{Z})$ has infinite index if and only if there exists a non-trivial element $\alpha' \in H_1(\partial\bar{N})$ such that $\varphi(\alpha') = 0$ for every $\varphi \in H$. Therefore, the index of $i^*(H^1(\bar{N}))$ in $H^1(\partial\bar{N})$ is infinite if and only if there exists a non-trivial element $\alpha' \in H_1(\partial\bar{N})$ such that

$$(12.1) \quad 0 = \langle i^*(\beta), \alpha' \rangle = \langle \beta, i_*(\alpha') \rangle \quad \text{for every } \beta \in H^1(\bar{N}).$$

An easy application of the Universal Coefficient Theorem shows that the Kronecker pairing between $H_1(\bar{N})$ and $H^1(\bar{N})$ induces an epimorphism $H^1(\bar{N}) \rightarrow \text{Hom}(H_1(\bar{N}), \mathbb{Z})$, so the condition described in Equation (12.1) is equivalent to the fact that $\varphi(i_*(\alpha')) = 0$ for every $\varphi \in \text{Hom}(H_1(\bar{N}), \mathbb{Z})$, whence to the fact that $i_*(\alpha')$ has finite order in $H_1(\bar{N})$.

We have thus shown that property (2) above is equivalent to the existence of a non-trivial element $\alpha' \in H_1(\partial\bar{N})$ such that $i_*(\alpha')$ has finite order in $H_1(\bar{N})$. Since $H_1(\partial\bar{N})$ is torsion-free, this last condition holds if and only if the kernel of i_* is non-trivial, concluding the proof. \square

Now the advantage in changing to a homological criterion is that it is easier to achieve geometrically. One needs to find examples of finite volume, non-compact, hyperbolic manifolds N having the property that they contain an embedded $S \hookrightarrow N$, where S is non-compact surface with finitely many cusps, and the embedding is proper. After truncation, this yields an element in $H_1(\partial\bar{N})$, namely the element corresponding to $\partial\bar{S} \hookrightarrow \partial\bar{N}$, having the property that $i_*([\partial\bar{S}]) = 0 \in H_1(\bar{N})$. Moreover, if S is suitably chosen one may also ensure that $[\partial\bar{S}] \neq 0$ in $H_1(\bar{N})$.

One approach to finding such examples would be to construct N so as to contain a properly embedded totally geodesic non-compact finite volume hyperbolic surface Σ . A forthcoming paper of McReynolds, Stover, and Reid [MRS] focuses on arithmetical constructions of such pairs (N, Σ) in all dimensions.

An alternate approach is to ignore the geometry and to try to argue purely topologically. Fixing a single boundary torus T inside one of these truncated hyperbolic

manifolds \bar{N} , we let x_1, \dots, x_{n-1} be a basis for the first cohomology $H^1(T) \cong \mathbb{Z}^{n-1}$. The following proposition was suggested to us by Juan Souto:

PROPOSITION 12.4. *Assume that the cohomology classes x_i for $1 \leq i \leq n-2$ have the property that $\langle x_i \rangle \cap i^*(H^1(\bar{N})) \neq \{0\}$. Then there exists an embedded smooth surface with boundary $(\Sigma, \partial\Sigma) \hookrightarrow (\bar{N}, \partial\bar{N})$, having the following properties:*

- i) $\Sigma \cap \partial\bar{N} = \partial\Sigma$ is entirely contained in the boundary component T , and*
- ii) the collection of curves $\partial\Sigma$ represent a non-zero class in $H_1(T)$ (and in particular, $\partial\Sigma \neq \emptyset$).*

PROOF. Since each $\langle x_i \rangle \cap i^*(H^1(\bar{N})) \neq \{0\}$, we can find non-zero integers r_1, \dots, r_{n-2} with the property that $r_i \cdot x_i \in i^*(H^1(\bar{N}))$ for $1 \leq i \leq n-2$. Let $y_i \in H^1(\bar{N})$ be chosen so that $i^*(y_i) = r_i \cdot x_i$. We will be considering elements in four (co)-homology groups, related via have the commutative diagram:

$$\begin{array}{ccc} H^1(\bar{N}) & \xrightarrow{i^*} & H^1(\partial\bar{N}) \\ \downarrow \cong & & \downarrow \cong \\ H_{n-1}(\bar{N}, \partial\bar{N}) & \xrightarrow{\partial} & H_{n-2}(\partial\bar{N}) \end{array}$$

where the vertical maps are isomorphisms given by Poincaré-Lefschetz duality, the top map is induced by inclusion, and the bottom map is the boundary map. We now proceed to use the cohomology classes y_i to construct the surface Σ .

First, recall that for a smooth k -manifold M (possibly with boundary), the Poincaré-Lefschetz dual of a 1-dimensional cohomology class $x \in H^1(M)$ has a simple geometric interpretation. One can think of the element x as a homotopy class of maps into the classifying space $K(\mathbb{Z}, 1) = S^1$, with the trivial element corresponding to a constant map. Fixing a reference point $p \in S^1$, we can find a smooth map f within the homotopy class with the property that f is transverse to p . Then $f^{-1}(p)$ defines a smooth submanifold, which represents the Poincaré-Lefschetz dual to x . This will represent a class in either $H_{k-1}(M)$ or in $H_{k-1}(M, \partial M)$, according to whether $\partial M = \emptyset$ or $\partial M \neq \emptyset$. For example, in the special case consisting of the trivial cohomology class, one can perturb the constant map to not contain p in the image, so that the dual class is represented by the “vacuous” submanifold.

Let us apply this procedure to each of the cohomology classes $y_i \in H^1(\bar{N})$, obtaining corresponding smooth maps $f_i : \bar{N} \rightarrow S^1$ transverse to p . Now the restriction of f_i to $\partial\bar{N}$ will yield the Poincaré-Lefschetz dual to the cohomology class $i^*(y_i) = r_i \cdot x_i \in H^1(\partial\bar{N})$. The cohomology $H^1(\partial\bar{N})$ decomposes as a direct sum of the cohomology of the individual boundary components, and by construction the class $i^*(y_i) = r_i \cdot x_i$ is purely supported on the $H^1(T)$ summand. Geometrically, this just says that the restriction of f_i to any of the remaining boundary components is homotopic to a point, which we can take to be distinct from p . Using a collared neighborhood of each of the boundary components, we can effect such a homotopy, allowing us to replace f_i by a homotopic map which has the additional property that T is the only boundary component of \bar{N} whose image intersects p .

Taking pre-images of p under these maps, we obtain a collection of $(n-1)$ -dimensional manifolds W_1, \dots, W_{n-2} representing the dual homology classes

in $H_{n-1}(\overline{N}, \partial\overline{N})$. Moreover, each W_i intersects $\partial\overline{N}$ in a collection of $(n-2)$ -dimensional submanifolds $\partial W_i \subset T$, which represent the duals to the cohomology classes $r_i \cdot x_i \in H^1(T)$. Perturbing the pairs $(W_i, \partial W_i) \subset (\overline{N}, T)$ slightly, we may assume they are all pairwise transverse. This in turn ensures that the intersection $\Sigma = \cap_{i=1}^{n-2} W_i$ is a smooth submanifold. Since Σ is the intersection of $n-2$ manifolds each of which has codimension one, we see that Σ has codimension $n-2$ in the n -dimensional manifold \overline{N} , i.e. Σ is a surface. Since T is the only boundary component which intersects any of the W_i , we have that $\partial\Sigma \subset T$ giving us (i).

So to conclude, we need to verify property (ii): that the family of curves defined by $\partial\Sigma$ represent a non-zero class in $H^1(T)$. But recall that $\partial\Sigma = \cap_{i=1}^{n-2} \partial W_i$, where each ∂W_i is an $(n-2)$ -dimensional submanifold of the $(n-1)$ -dimensional torus T , representing the Poincaré dual to the cohomology class $i^*(y_i) = r_i \cdot x_i \in H^1(T)$. Under Poincaré duality, the geometric intersection of cycles corresponds to the cup product of the dual cocycles. As such, the collection of curves $\partial\Sigma$ represents the Poincaré dual of the cup product

$$\cup_{i=1}^{n-2} (r_i \cdot x_i) = \left(\prod r_i \right) \cdot \left(\cup_{i=1}^{n-2} x_i \right) \in H^{n-2}(T) \cong \mathbb{Z}^{n-1}.$$

We know that the cohomology ring $H^*(T)$ is an exterior algebra over the x_i , hence the cup product $\cup_{i=1}^{n-2} x_i$ is non-zero. Since the coefficient $\prod r_i$ is a non-zero integer, the Poincaré dual of $[\partial\Sigma] \in H_1(T)$ is non-trivial. This implies that the homology class $[\partial\Sigma]$ is likewise non-zero, establishing (ii), and concluding the proof of the Proposition. \square

COROLLARY 12.5. *The map $i_* : H_1(\partial\overline{N}) \rightarrow H_1(\overline{N})$ is not injective.*

PROOF. Fix a boundary component T of \overline{N} , and choose a basis x_1, \dots, x_{n-1} for the first cohomology $H^1(T) \cong \mathbb{Z}^{n-1}$. If any of the elements x_1, \dots, x_{n-2} have the property that $\langle x_i \rangle \cap i^*(H^1(\overline{N})) = \{0\}$, then we are done by Lemma 12.3. So we can assume that $\langle x_i \rangle \cap i^*(H^1(\overline{N})) \neq \{0\}$ for each $1 \leq i \leq n-2$, allowing us to apply Proposition 12.4, whence the conclusion again. \square

Putting together Proposition 12.2, Lemma 12.3 and Corollary 12.5, we can now establish:

THEOREM 12.6. *Let N be any finite volume, non-compact, hyperbolic manifold, with all cusps diffeomorphic to a torus times $[0, \infty)$, and let \overline{N} be the compact manifold obtained by “truncating the cusps”. Then one can find a graph manifold, arising as a principal S^1 -bundle over the double $D\overline{N}$, which does not support a locally CAT(0) metric.*

To conclude, we recall that there exist examples, in all dimensions ≥ 3 , of non-compact finite volume hyperbolic manifolds with toric cusps (see [MRS]). From Theorem 12.6, we immediately deduce:

COROLLARY 12.7. *There are examples, in all dimensions ≥ 4 , of principal S^1 -bundles which are graph manifolds, but do **not** support any locally CAT(0) metric.*

12.2. Irreducible examples

Keeping the notation from Section 2.1, let N be a complete finite-volume hyperbolic manifold with toric cusps, and set $V = \overline{N} \times S^1$, where \overline{N} is as usual the

natural compactification of N . We denote by n the dimension of V . In this Subsection we show that one may always choose affine gluing maps between the boundary components of two copies of V in such a way that the resulting graph manifold is irreducible and does not support any locally CAT(0) metric. As a consequence, irreducible graph manifolds which do not support any locally CAT(0) metric exist in every dimension ≥ 4 .

Let T_1^*, \dots, T_r^* be the boundary components of V . We denote by V^+, V^- two copies of V , and by T_i^+ (resp. by T_i^-) the boundary component of V^+ (resp. of V^-) corresponding to T_i^* , $i = 1, \dots, r$. For every $i = 1, \dots, r$ we fix an affine diffeomorphism $\psi_i: T_i^+ \rightarrow T_i^-$, we denote by M the graph manifold obtained by gluing V^+ and V^- along the ψ_i , and by $T_i \subseteq M$ the torus corresponding to $T_i^+ \subseteq \partial V^+$ and $T_i^- \subseteq \partial V^-$.

Let us now suppose that M admits a locally CAT(0) metric. Following [Le], we briefly describe the Euclidean scalar product induced by such a metric on each $H_1(T_i)$, $i = 1, \dots, r$. We endow \widetilde{M} with the CAT(0) metric induced by the covering map, and denote by A_i the subgroup of the group $\Gamma = \text{Aut}(\widetilde{M}) \cong \pi_1(M)$ corresponding to (a representative of the conjugacy class of) $\pi_1(T_i)$. Let us briefly recall some definitions and results from [BrHa, Chapter II.6]. For every $\gamma \in \Gamma$ the *translation length* of γ is defined by setting

$$\tau(\gamma) = \inf\{d(x, \gamma(x)) \mid x \in \widetilde{M}\}.$$

We also set

$$\text{Min}(\gamma) = \{x \in \widetilde{M} \mid d(x, \gamma(x)) = \tau(\gamma)\} \subseteq \widetilde{M}.$$

Since $\widetilde{M}/\Gamma = M$ is compact, the isometry γ is semisimple, and $\text{Min}(\gamma)$ is non-empty. If H is a subgroup of Γ we finally set $\text{Min}(H) = \bigcap_{\gamma \in H} \text{Min}(\gamma)$.

Since $A_i \cong \mathbb{Z}^{n-1}$, by the Flat Torus Theorem the subset $\text{Min}(A_i)$ splits as a metric product $\text{Min}(A_i) = Y_i \times E^{n-1}$, where E^k is the Euclidean k -dimensional space (see *e.g.* [BrHa, Chapter II.7]). Moreover, A_i leaves $\text{Min}(A_i)$ invariant, and the action of every $a \in A_i$ on $\text{Min}(A_i)$ splits as the product of the identity on Y_i and a non-trivial translation $v \mapsto v + v_a$ on E^{n-1} . If l_1, l_2 are elements of $H_1(T_i)$ we set

$$\langle l_1, l_2 \rangle_i = \langle v_{a_1}, v_{a_2} \rangle,$$

where a_j is the element of $A_i \cong \pi_1(T_i) \cong H_1(T_i)$ corresponding to l_j , and $\langle \cdot, \cdot \rangle_i$ denotes the standard scalar product of E^{n-1} . It is readily seen that $\langle \cdot, \cdot \rangle_i$ is indeed well-defined. Moreover, the norm $\|l\|_i = \sqrt{\langle l, l \rangle_i}$ of any element $l \in H_1(T_i)$ coincides with the translation length of the corresponding element $a \in A_i < \Gamma$, so if $l_1, l_2 \in H_1(T_i)$ correspond to the elements $a_1, a_2 \in A_i$ we have

$$2\langle l_1, l_2 \rangle_i = \tau(a_1 \circ a_2)^2 - \tau(a_1)^2 - \tau(a_2)^2.$$

Let us now denote by $f^\pm \in H_1(V^\pm)$ the class represented by the fiber of V^\pm , *i.e.* the element of $H_1(V^\pm) = H_1(N) \oplus H_1(S^1)$ corresponding to the positive generator of $H_1(S^1)$. If $i_*^\pm: \bigoplus_{i=1}^r H_1(T_i) \rightarrow H_1(V^\pm)$ is the map induced by the inclusion $i^\pm: \bigsqcup_{i=1}^r T_i \rightarrow V^\pm$ (where we consider V^\pm as a subset of M), then for every $i = 1, \dots, r$ there exists a unique element $f_i^\pm \in H_1(T_i)$ such that $i_*^\pm(f_i^\pm) = f^\pm$. Observe that our definitions imply that M is irreducible if and only if $f_i^+ \neq \pm f_i^-$ for every $i = 1, \dots, r$. Lemma 12.8 and Proposition 12.9 below are inspired by the proof of [KaLe2, Theorem 3.7]:

LEMMA 12.8. *For every $i = 1, \dots, r$ let b_i be an element of $H_1(T_i)$ such that*

$$i_*^\pm(b_1 + \dots + b_r) = 0 .$$

Then

$$\sum_{i=1}^r \langle b_i, f_i^\pm \rangle_i = 0 .$$

PROOF. Let us fix a representative Γ^\pm of the conjugacy class of the subgroup $\pi_1(V^\pm)$ of $\pi_1(M) \cong \Gamma$. We also choose the subgroups A_i corresponding to the tori T_i in such a way that $A_i < \Gamma^\pm$ for every $i = 1, \dots, r$. Let $\phi \in \Gamma^\pm$ be the element corresponding to $(\text{Id}, 1)$ under the identification

$$\Gamma^\pm = \pi_1(V^\pm) = \pi_1(\overline{N}) \times \pi_1(S^1) = \pi_1(\overline{N}) \times \mathbb{Z} .$$

By construction we have $\phi \in \bigcap_{i=1}^r A_i \subseteq \Gamma^\pm$, and the image of ϕ under the Hurewicz homomorphism $\pi_1(V^\pm) \rightarrow H_1(V^\pm)$ coincides with f^\pm .

Since ϕ lies in the center of Γ^\pm the set $\text{Min}(\phi) \subseteq \widetilde{M}$ is Γ^\pm -invariant. Moreover, the action of Γ^\pm preserves the isometric splitting $\text{Min}(\phi) = W \times E^1$, so the identification of Γ with the group of covering automorphisms of \widetilde{M} induces a representation $\rho: \Gamma^\pm \rightarrow \text{Isom}(W) \times \text{Isom}(E^1)$. If $\rho_0: \Gamma^\pm \rightarrow \text{Isom}(W)$, $\rho_1: \Gamma^\pm \rightarrow \text{Isom}(E^1)$ are the components of ρ , then $\rho_0(\phi)$ is the identity of W , while $\rho_1(\phi)$ is a non-trivial translation. As a consequence, since for every $\gamma \in \Gamma^\pm$ the isometries $\rho_1(\gamma)$ and $\rho_1(\phi)$ commute, the representation ρ_1 takes values in the abelian group of translations of E^1 , which can be canonically identified with \mathbb{R} . Therefore, by Hurewicz's Theorem the homomorphism ρ_1 must factor through $H_1(V^\pm)$, thus defining a homomorphism $\bar{\rho}_1: H_1(V^\pm) \rightarrow \mathbb{R}$.

Let us now observe that, since $\phi \in A_i$, we have $\text{Min}(A_i) \subseteq \text{Min}(\phi) = W \times E^1$, so in order to compute the translation length of elements of A_i it is sufficient to consider their action on $W \times E^1$. Therefore, for every $a \in A_i$ we have $\tau(a)^2 = \tau_W(\rho_0(a))^2 + \rho_1(a)^2$, where we denote by τ_W the translation length of elements of $\text{Isom}(W)$, and we recall that we are identifying the group of translations of E^1 with \mathbb{R} . We now let $\beta_i \in A_i$ be a representative of $b_i \in H_1(T_i)$, and proceed to evaluate the scalar product $\langle b_i, f_i^\pm \rangle_i$. We know that:

$$2\langle b_i, f_i^\pm \rangle_i = \tau(\phi \circ \beta_i)^2 - \tau(\phi)^2 - \tau(\beta_i)^2 .$$

Considering the terms on the right hand side, we recall that $\phi \in A_i$ is a representative of $f_i^\pm \in H_1(T_i)$, and hence we have $\tau(\phi)^2 = \rho_1(\phi)^2$. Using the product structure on $W \times E^1$, the remaining two terms are $\tau(\beta_i)^2 = \tau_W(\rho_0(\beta_i))^2 + \rho_1(\beta_i)^2$, and $\tau(\phi \circ \beta_i)^2 = \tau_W(\rho_0(\beta_i))^2 + (\rho_1(\phi) + \rho_1(\beta_i))^2$. Substituting these into the expression and simplifying, we obtain that

$$2\langle b_i, f_i^\pm \rangle_i = 2\rho_1(\phi)\rho_1(\beta_i) = 2\rho_1(\phi)\bar{\rho}_1(i_*^\pm(b_i)) .$$

Summing over all i , we deduce that

$$\sum_{i=1}^r \langle b_i, f_i^\pm \rangle = \rho_1(\phi) \cdot \sum_{i=1}^r \bar{\rho}_1(i_*^\pm(b_i)) = \rho_1(\phi) \cdot \bar{\rho}_1 \left(i_*^\pm \left(\sum_{i=1}^r b_i \right) \right) = 0 ,$$

whence the conclusion. \square

PROPOSITION 12.9. *There exists a choice for the gluing maps $\psi_i: T_i^+ \rightarrow T_i^-$ such that the following conditions hold:*

- (1) the graph manifold M obtained by gluing V^+ and V^- along the ψ_i 's is irreducible;
- (2) M does not admit any locally CAT(0) metric.

PROOF. Let Y_1, \dots, Y_r be the boundary components of \bar{N} . By Corollary 12.5, there exist elements $b'_i \in H_1(Y_i)$, $i = 1, \dots, r$, such that $0 \neq b'_1 + \dots + b'_r \in H_1(Y_1) \oplus \dots \oplus H_1(Y_r) = H_1(\partial\bar{N})$, and $i_*(b'_1 + \dots + b'_r) = 0$ in $H_1(\bar{N})$, where i_* is induced by the inclusion $\partial\bar{N} \hookrightarrow \bar{N}$. Recall that $V^\pm = \bar{N} \times S^1$, so that we have natural identifications $T_i^\pm = Y_i \times S^1$ and $H_1(T_i^\pm) = H_1(Y_i \times S^1) \cong H_1(Y_i) \oplus H_1(S^1)$, $i = 1, \dots, r$. Under this identifications, every affine diffeomorphism $\psi_i: T_i^+ \rightarrow T_i^-$ induces an isomorphism

$$(\psi_i)_*: H_1(Y_i) \oplus H_1(S^1) \rightarrow H_1(Y_i) \oplus H_1(S^1).$$

Let us denote by λ the positive generator of $H_1(S^1)$. For every $i = 1, \dots, r$, we choose the diffeomorphism $\psi_i: T_i^+ \rightarrow T_i^-$ as follows. Let $I = \{i \mid b'_i \neq 0\} \subset \{1, \dots, r\}$, and observe that our assumptions ensure that I is non-empty. If $i \notin I$, we only ask that $(\psi_i)_*(0, \lambda) \neq (0, \pm\lambda)$, so that $f_i^+ \neq \pm f_i^-$ in $H_1(T_i)$, where we recall that T_i is the toric hypersurface corresponding to T_i^+ and T_i^- in the resulting graph manifold M . If $i \in I$, we choose a positive integer n_i and we let ψ_i be an affine diffeomorphism such that $(\psi_i)_*(v, 0) = (v, 0)$ for every $v \in H^1(Y_i)$ and $(\psi_i)_*(0, \lambda) = (n_i b'_i, \lambda)$.

Our choices imply that for every $i \in I$ there exists a well-defined element $b_i \in H_1(T_i)$ corresponding to the elements $(b_i, 0) \in H_1(T_i^+)$ and $(b_i, 0) = (\psi_i)_*(b_i, 0) \in H_1(T_i^-)$. We also set $b_i = 0 \in H_1(T_i)$ for every $i \notin I$, and observe that $f_i^+ = f_i^- + n_i b_i$ for every $i \in I$.

Let M be the graph manifold obtained by gluing V^+ and V^- along the ψ_i 's. By construction, M is irreducible. Let us suppose by contradiction that M admits a locally CAT(0) metric. It is readily seen that $i_*^\pm(\sum_{i=1}^r b_i) = 0$, so Lemma 12.8 implies that

$$\begin{aligned} 0 &= \sum_{i=1}^r \langle f_i^+, b_i \rangle_i = \sum_{i=1}^r \langle f_i^- + n_i b_i, b_i \rangle_i \\ &= \sum_{i=1}^r \langle f_i^-, b_i \rangle_i + \sum_{i=1}^r n_i \langle b_i, b_i \rangle_i \\ &= \sum_{i=1}^r n_i \|b_i\|_i^2, \end{aligned}$$

a contradiction since $n_i > 0$ and $b_i \neq 0$ for every $i \in I$. We have thus shown that M cannot support any locally CAT(0) metric, and this concludes the proof. \square

COROLLARY 12.10. *For every integer $n \geq 4$, there exist infinitely many closed irreducible graph n -manifolds which do **not** support any locally CAT(0) metric.*

PROOF. Let us fix an integer $m \geq 3$. It is proved in [MRS] that there exist infinitely many complete finite-volume hyperbolic m -manifolds with toric cusps. If N is any such manifold, Proposition 12.9 shows that there exists an irreducible graph manifold M which does not support any locally CAT(0) metric and decomposes as the union of two pieces V^+ and V^- , each of which is diffeomorphic to $\bar{N} \times S^1$.

In order to conclude it is sufficient to show that the diffeomorphism type of M completely determines the hyperbolic manifold N , so that the infinite family of

hyperbolic manifolds provided by [MRS] gives rise the infinite family of desired examples. However, Theorem 0.3 implies that the diffeomorphism type of M determines the isomorphism type of the fundamental group of V^\pm . Since $\pi_1(N)$ is equal to the quotient of $\pi_1(V^\pm)$ by its center (see Remark 2.9), the conclusion follows by Mostow rigidity. \square

REMARK 12.11. Even when starting with a fixed pair of pieces, one can still obtain an infinite family of irreducible graph manifolds which do not support any locally CAT(0) metric. For example, let N be a hyperbolic knot complement in S^3 , set $V^+ = V^- = \overline{N} \times S^1$ and denote by T^+ (resp. T^-) the unique boundary component of V^+ (resp. of V^-). The boundary of a Seifert surface for K defines an element $b' \in H_1(\partial\overline{N})$ which bounds in \overline{N} , whence an element $b \in H_1(T^\pm)$ such that $i_*(b) = 0 \in H_1(V^\pm)$. Let $M(n)$ be the irreducible graph manifold obtained by gluing the base of V^+ to the base of V^- via the identity of $\partial\overline{N}$, and by gluing the fibers of V^+ and V^- in such a way that $f^+ = f^- + nb$ in $H_1(T)$, where T is the toric hypersurface in $M(n)$ corresponding to T^+ and T^- . As described in the proof of Proposition 12.9, for every positive integer n the manifold $M(n)$ does not admit any locally CAT(0) metric. Moreover, as explained in Remark 8.11, the proof of Theorem 8.9 can be adapted to show that among the fundamental groups of the $M(n)$'s, there are infinitely many non-isomorphic groups.

REMARK 12.12. Let N be a complete finite-volume hyperbolic manifold with toric cusps. We have proved in Proposition 12.9 that there exist “twisted doubles” of $\overline{N} \times S^1$ which provide examples of *closed* irreducible graph manifolds not admitting any locally CAT(0) metric. However, in principle one can use a similar construction to also get examples with non-empty boundary.

Indeed, if $T_1 \cup \dots \cup T_k \subseteq \partial\overline{N} \times S^1$ is a family of boundary tori such that the map $i_* : H_1(T_1 \cup \dots \cup T_k) \rightarrow H_1(\overline{N} \times S^1)$ is not injective, then the proof of Proposition 12.9 shows that the obstruction to putting a global nonpositively curved metric on such twisted doubles is concentrated near the gluing tori T_1, \dots, T_k . In other words, if $\partial(\overline{N} \times S^1)$ contains some boundary component other than T_1, \dots, T_k , we can easily construct irreducible graph manifolds just by gluing two copies of $\overline{N} \times S^1$ along the corresponding copies of T_1, \dots, T_k , thus obtaining examples of irreducible graph manifolds, with non-empty boundary, and which do not support any locally CAT(0) metric.

Allowing surface pieces

13.1. Pieces with a surface base

A restriction that we have imposed on our graph manifolds is that all pieces have a base which is hyperbolic of *dimension* ≥ 3 . The reason for this restriction is obvious: hyperbolic manifolds of dimension ≥ 3 exhibit a lot more rigidity than surface groups. In this section we point out how to extend our results to a more general case, namely we now allow surfaces with boundary as bases of pieces.

DEFINITION 13.1. For $n \geq 3$, an *extended graph n -manifold* is a manifold built up from pieces as in the definition of graph manifold as well as *surface pieces*, that is manifolds of the form $\Sigma \times T^{n-2}$ with Σ non-compact, finite volume, hyperbolic surface. Also, we require that each gluing does not identify the fibers in adjacent surface pieces.

Let us briefly comment about the last requirement described in the above Definition. If we allowed gluings which identify the fibers of adjacent surface pieces, then the resulting decomposition into pieces of our extended graph manifold would no longer be canonical. Indeed, within a surface piece $\Sigma \times T^{n-2}$, we can take any non-peripheral simple closed curve $\gamma \hookrightarrow \Sigma$ in the base surface, and cut the piece open along $\gamma \times T^{n-2}$. This allows us to break up the original piece $\Sigma \times T^{n-2}$ into pieces $(\Sigma \setminus \gamma) \times T^{n-2}$ (which will either be two pieces, or a single “simpler” piece, according to whether γ separates or not). Our additional requirement avoids this possibility. Note however that if one has adjacent surface pieces with the property that the gluing map matches up their fibers exactly, then it is not possible to conclude that the two surface pieces can be combined into a single surface piece (the resulting manifold could be a non-trivial S^1 -fiber bundle over a surface rather than just a product).

Some of our results extend to the broader class of extended graph manifolds for free. For instance, the argument for topological rigidity given in Chapter 3 goes through unchanged (see the discussion in Remark 3.3). Amongst the algebraic properties of $\pi_1(M)$ that we discuss in Chapter 6, the following extend verbatim to the more general context:

- (i) that there are no Kazhdan subgroups,
- (ii) that the Tits alternative holds,
- (iii) that they have uniform exponential growth.

In order to extend our rigidity results to cover (certain) extended graph manifolds, the key again lies in showing that fundamental groups of walls are preserved. Let M be an extended graph n -manifold, with fundamental group Γ and Bass-Serre tree T . As a first step towards our goal, we have the easy:

LEMMA 13.2. *If $H \leq \Gamma$ is a subgroup isomorphic to \mathbb{Z}^{n-1} , then either:*

- (i) H stabilizes an edge in T , or
- (ii) H stabilizes a vertex v in T , and v corresponds to a surface piece in M .

PROOF. The proof is a minor variation of Lemma 4.1. $H \cong \mathbb{Z}^{n-1}$ is an abelian group acting on a tree T , so either it fixes a vertex, or it stabilizes a geodesic γ (on which it acts via translations). If H stabilizes a vertex v , then either v corresponds to a surface piece, and we obtain (ii), or v corresponds to a piece with a hyperbolic base of dimension ≥ 3 , in which case H must stabilize an edge incident to v , giving us (i). So to conclude, we just need to argue that H cannot stabilize any geodesic $\gamma \subset T$.

The stabilizer $Stab(\gamma)$ of any geodesic has a subgroup $Fix(\gamma)$, with quotient isomorphic to either 1, $\mathbb{Z}/2$, \mathbb{Z} , or \mathbb{D}_∞ . So if $H = Stab(\gamma)$, then the subgroup $Fix(\gamma) \leq H \cong \mathbb{Z}^{n-1}$ is abstractly isomorphic to either (i) \mathbb{Z}^{n-1} or (ii) \mathbb{Z}^{n-2} . For the action on the Bass-Serre tree, we know that the subgroup which fixes a pair of adjacent edges, when thought of as a subgroup of the common vertex group, is contained in the corresponding fiber subgroup (see the trichotomy mentioned in Lemma 8.1). Since these fiber subgroups have rank $\leq n-2$, we see that (i) cannot occur.

To see that (ii) cannot occur, we note that this would force *all* vertices on the geodesic γ to correspond to surface pieces. But we assumed that surface pieces have fiber subgroups whose intersection has rank $\leq n-3$. Since H would have to be contained in this intersection, we again obtain a contradiction. This rules out case (ii), concluding the proof of the Lemma. \square

The lemma above is not enough to establish a full characterization of fundamental groups of walls. In the next section we establish the desired characterization.

13.2. Characterizing the wall subgroups

Consider the collection $F(M)$ of maximal subgroups of Γ which are isomorphic to \mathbb{Z}^{n-1} . We want to characterize the fundamental groups of walls in $F(M)$. We will say that $H \in F(M)$ is adjacent to a piece if it is contained in (a conjugate of) the fundamental group of a boundary component of the piece.

Notice that all $H \in F(M)$ contained in the stabilizer of the same vertex corresponding to a surface piece contain a specified subgroup isomorphic to \mathbb{Z}^{n-2} (namely, the corresponding fiber subgroup).

THEOREM 13.3 (Characterizing walls). *For each $H \in F(M)$,*

- (1) H is adjacent to distinct non-surface pieces if and only if for each $H' \in F(M)$ with $H' \neq H$ we have that $H \cap H'$ has rank less than $n-2$.
- (2) H is adjacent to distinct surface pieces if and only if there exist $K, K' \in F(M)$ such that $K \cap H, K' \cap H$ have rank $n-2$, but $K \cap K'$ has rank strictly less than $n-2$.
- (3) H is adjacent to a surface piece and to a non-surface piece if and only if the following holds. There exists $K \in F(M)$ such that $L = K \cap H$ has rank $n-2$. Also, there exists $H' \in F(M)$ with the property that for each D , and each finite collection $K_1, \dots, K_r \in F(M) \setminus \{H, H'\}$, there exists a path connecting H' to H and avoiding the set $N_D(L) \cup N_D(K_1) \cup \dots \cup N_D(K_r)$ (in the Cayley graph of Γ).

PROOF. Items (1) and (2) are easy, so let us focus on (3). Recall that every chamber of \widetilde{M} is endowed with an intrinsic CAT(0) path metric. Therefore, in view of Lemma 13.2, elements of $F(M)$ correspond geometrically to codimension one flats contained within a single chamber of \widetilde{M} (in case (i), the flat is not unique, as it can be chosen in either adjacent chamber - but the two flats are at a bounded distance apart). Along with the Milnor-Svarc Lemma and Lemma 2.16 (which ensures that our arguments work also in the case when chambers are not quasi-isometrically embedded in \widetilde{M}), this allows us to translate condition (3) into equivalent geometric statements concerning corresponding configurations of flats inside \widetilde{M} .

(\Rightarrow) We start by showing the “only if” part. The existence of K is clear. In order to verify the second property, it is enough to show the following geometric statement. We are given a wall W in \widetilde{M} adjacent to two chambers C_s, C_{ns} (where C_s covers a surface piece and C_{ns} covers a non-surface piece), a fiber $F \subset W$ of the chamber C_s , and a distinct wall W' adjacent to C_{ns} . Then for any prescribed finite collection of flats V_1, \dots, V_k (each $\neq W, W'$) and any prescribed $D \geq 0$, we can find a path β in C_{ns} connecting W to W' which avoids the set $N_D(F) \cup N_D(V_1) \cup \dots \cup N_D(V_k)$.

First, we note that it is enough to show this property in the special case where the flats V_1, \dots, V_k are chosen to lie within C_{ns} . Indeed, if some V_i is *not* within C_{ns} , then we can find a corresponding boundary wall W_i of C_{ns} with the property that any path from V_i to C_{ns} intersects W_i . Then from the triangle inequality we have that $N_D(V_i) \cap C_{ns} \subset N_D(W_i) \cap C_{ns}$. Replacing V_i by the corresponding W_i , we see that any path in C_{ns} avoiding the D -neighborhood of this new finite family of flats automatically avoids the D -neighborhood of the original family of flats. After performing such a replacement for each V_i which is not in C_{ns} , we have reduced to the case where all the flats are in C_{ns} .

Next, we observe that it is enough to verify that this property holds in C_{ns} endowed with its path metric (as mentioned above, this reduction does not require C_{ns} to be quasi-isometrically embedded, see Lemma 2.16). The chamber C_{ns} splits as a product $\widehat{\mathbb{H}}^k \times \mathbb{R}^{n-k}$, where $\widehat{\mathbb{H}}^k$ ($k \geq 3$) is a copy of hyperbolic space with a suitable equivariant family of horoballs removed. The walls W, W' decompose as products $B \times \mathbb{R}^{n-k}, B' \times \mathbb{R}^{n-k}$, where we have B, B' are boundary components of $\widehat{\mathbb{H}}^k$. Similarly, each flat V_i decomposes as a product $B_i \times \mathbb{R}^{n-k}$, where again, each B_i is a boundary component of $\widehat{\mathbb{H}}^k$. Recall that, within the wall $W = B \times \mathbb{R}^{n-k}$, we also have a copy of the fiber F of the adjacent surface chamber C_s . F defines a codimension one flat subspace of W , and we now have two possibilities to consider: either (a) $F \cap (B \times \{0\}) = B \times \{0\}$, or (b) $F \cap (B \times \{0\})$ defines a proper subspace of $B \times \{0\}$.

In case (a), the fact that F has codimension one in the wall $B \times \mathbb{R}^{n-k}$ and contains $B \times \{0\}$ implies that, for any $D \geq 0$, one can find a point v in the fiber \mathbb{R}^{n-k} so that the corresponding slice $\widehat{\mathbb{H}}^k \times \{v\}$ has empty intersection with $N_D(F)$. We also see that each of the sets $N_D(V_i)$ intersects that slice in the set $N_D(B_i) \times \{v\}$. Focusing on that slice, we now claim that we can connect a point in $B \times \{v\}$ to a point in $B' \times \{v\}$ within that slice, while avoiding the sets $N_D(B_i) \times \{v\}$. Indeed, in the proof of Proposition 4.3, we showed that, in any such $\widehat{\mathbb{H}}^k$ (with $k \geq 3$), and for any $D > 0$, we can connect any given pair of boundary components by a path which stays at least distance D away from any single other boundary component.

We now want to show the analogous result where the path stays at least distance D from a fixed *finite collection* of other boundary components.

To do this, we fix an isometric embedding of $\widehat{\mathbb{H}}^k$ into the upper half-space model for \mathbb{H}^k , chosen so that one of the deleted horoballs (distinct from those we are interested in) is centered at the point at infinity. With respect to the Euclidean metric on the upper half-space $\mathbb{R}^{k-1} \times (0, \infty)$, the remaining horoballs which have been removed are a family of pairwise disjoint Euclidean balls $\{O_\alpha\}$, each of which is tangent to the hyperplane $\mathbb{R}^{k-1} \times \{0\}$. Moreover, this family is invariant under an action of $H := \mathbb{Z}^{k-1}$ (the subgroup stabilizing the point at infinity) acting cocompactly via suitable translations on the \mathbb{R}^{k-1} factor. We will consider various horizontal slices $Y_h := \mathbb{R}^{k-1} \times \{h\}$, where $h > 0$ is some small positive number. The complement of the open horoballs $\{O_\alpha\}$ in Y_h will be denoted \widehat{Y}_h . The fact that the horoballs are pairwise disjoint, equivariant under the H -action, and have pairwise (hyperbolic) distance bounded away from zero implies the following basic facts:

- for each $h > 0$, there are only finitely many H -orbits of horoballs which intersect the prescribed horizontal slice Y_h .
- there is a function $r(h)$, with the property that the intersection of any horoball with the slice Y_h is a ball of (Euclidean) radius $\leq r(h)$, and the function has the property that $\lim_{h \rightarrow 0} r(h) = 0$.
- there is a positive function $d(h)$, with the property that inside \widehat{Y}_h , the (Euclidean) distance between any two distinct boundary components is $\geq d(h)$.

As a consequence, each \widehat{Y}_h is obtained from the corresponding copy Y_h of the Euclidean plane \mathbb{R}^{k-1} ($k \geq 3$) by removing the interior of a countably infinite (once h is small enough) family of embedded Euclidean balls, whose pairwise distance is bounded away from zero. This implies that all the sets \widehat{Y}_h (h small) are pairwise homeomorphic to each other. In fact, we have that each such \widehat{Y}_h is a non-compact, connected manifold with boundary, and the boundary components correspond precisely to the intersection of Y_h with the boundary spheres for the (Euclidean) balls $\{O_\alpha\}$. The argument for Proposition 4.3 relied on the fact that \widehat{Y}_h is path connected. Note that removing boundary components from a connected manifold with boundary doesn't affect the homotopy type of the space, hence the $D = 0$ case of our claim already follows.

For the stronger result we require, we note that fixing a $D > 0$ and one of the horoballs in our collection, the D -neighborhood is a horoball centered at the same point, of slightly larger radius. In terms of the upper half-space model, we replace the corresponding (Euclidean) ball O tangent to $\mathbb{R}^{k-1} \times \{0\}$ at the point p by a slightly larger (Euclidean) ball \widehat{O} tangent to $\mathbb{R}^{k-1} \times \{0\}$ at the same point p . At the level of one of the slices Y_h , the original horoball intersected Y_h in a ball $Y_h \cap O$ centered at some point $q \in Y_h$ of (Euclidean) radius $r_O(h)$, while the slightly enlarged horoball will intersect Y_h in a ball $Y_h \cap \widehat{O}$ centered at the same point q , but of slightly larger radius $r_O(h) + \delta_{O,D}(h)$. For the fixed $D > 0$ and chosen horoball, the positive function $\delta_{O,D}(h)$ satisfies $\lim_{h \rightarrow 0} \delta_{O,D}(h) = 0$.

Having set up the terminology and notation, we now return to our original goal. The finite collection of horoballs determine a finite collection of corresponding (Euclidean) balls O_1, \dots, O_k in the upper half-space, each of which has corresponding point of tangency p_1, \dots, p_k to $\mathbb{R}^{k-1} \times \{0\}$. The two distinguished horoballs we are trying to connect determine a pair of (Euclidean) balls O, O' , with corresponding points of tangency p, p' on $\mathbb{R}^{k-1} \times \{0\}$. Let Δ be a real number which is smaller than all (Euclidean) distances occurring between pairs of distinct points in the set $\{p, p', p_1, \dots, p_k\}$. Since $\lim_{h \rightarrow 0} r(h) = 0$, and each $\delta_{O_i, D}(h)$ satisfies $\lim_{h \rightarrow 0} \delta_{O_i, D}(h) = 0$, we can choose h small enough to ensure that $3r(h) + 2 \max_i \delta_{O_i, D}(h) < \Delta$. We make the:

Claim: The space $\bar{Y}_h := \hat{Y}_h \setminus \cup \bar{O}_i$, where the \bar{O}_i are the slightly larger horoballs centered at the points p_i , is still a manifold with boundary, homeomorphic to the original \hat{Y}_h , and contains amongst its boundary components the spheres $\partial O \cap Y_h$ and $\partial O' \cap Y_h$.

Assuming this, the result would immediately follow, as we could connect the two boundary components $\partial O \cap Y_h, \partial O' \cap Y_h$, via a path passing through the interior of the manifold with boundary \bar{Y}_h . Such a path would, by construction, stay at distance $\geq D$ from the finitely many horoballs O_1, \dots, O_k .

To see the Claim, we first observe that each $\bar{O}_i \cap Y_h$ intersects only finitely many boundary components of \bar{Y}_h (by finiteness of the number of H -orbits of these boundary components). These finitely many boundary components came about by removing finitely many open balls $\{U_1^i, \dots, U_{s_i}^i\}$ from Y_h . Rather than first removing the U_s^i (to get \hat{Y}_h) and then removing the set \bar{O}_i , we can instead just directly remove the set $U_i := \bar{O}_i \cup \bigcup_{s=1}^{s_i} U_s^i$ from Y_h . Note that each of the sets U_i is homeomorphic to an open ball, as it is the union of finitely many pairwise disjoint open Euclidean balls which intersect a fixed open Euclidean ball. The boundary of the set U_i is a tame embedding of a codimension one sphere. At the cost of perturbing h slightly, we can also assume that the distance from U_i to any of the remaining open balls is strictly positive. The sets U_1, \dots, U_k are pairwise disjoint, since a non-empty intersection $U_i \cap U_j \neq \emptyset$ could only occur if there is an open ball U_s^i intersecting both \bar{O}_i and \bar{O}_j . The triangle inequality would then imply that the Euclidean distance between p_i, p_j is bounded above by:

$$d(p_i, p_j) \leq (r(h) + \delta_{O_i, D}(h)) + r(h) + (r(h) + \delta_{O_j, D}(h)) < \Delta$$

contradicting the choice of Δ . Putting this together, we see that \bar{Y}_h is also obtained from Y_h by removing the interiors of a countably infinite collection of pairwise disjoint, tamely embedded balls (the finitely many U_i , along with all the other Euclidean balls that didn't intersect any of the \bar{O}_i), whose pairwise distance is bounded away from zero. But this implies that \bar{Y}_h is again homeomorphic to \hat{Y}_h , and in particular, is a manifold with boundary. Finally, the set $\partial O \cap Y_h$ is a boundary component of \bar{Y}_h , for otherwise the set $O \cap Y_h$ would have to be one of the U_s^i . But then the triangle inequality would allow us to estimate the Euclidean distance between p, p_i from above:

$$d(p, p_i) \leq (r(h) + \delta_{O_i, D}(h)) + r(h) < \Delta$$

which again contradicts the choice of Δ . An identical argument gives that $\partial O' \cap Y_h$ is a boundary component of \bar{Y}_h . This completes the construction in case (a).

In case (b), we focus instead on the slice $\widehat{\mathbb{H}}^k \times \{0\}$. The set F intersects $B \times \{0\}$ in a codimension one affine subspace, call it F_0 , and the subset $N_D(F)$ intersects the slice in the set $N_{D'}(F_0) \times \{0\}$ (where D' could be different than D , depending on the “slope” of F). As before, each $N_D(V_i)$ intersects the slice in the set $N_D(B_i) \times \{0\}$. So it is enough to show that, within the $\widehat{\mathbb{H}}^k$, given a collection of pairwise distinct boundary components B, B', B_1, \dots, B_k , a codimension one affine subspace $F_0 \subset B$, and a pair of real numbers $D, D' > 0$, one can always find a path joining a point in B to a point in B' , while staying at least D' away from F_0 , and at least D away from each B_i .

But this can easily be resolved using the arguments from case (a). Fix an identification of $\widehat{\mathbb{H}}^k$ with the upper half-plane model, so that the horoball corresponding to the boundary component B is centered at the point at infinity. The boundary component B now corresponds to a slice Y_H at some fixed (large) height H . The subspace F_0 is a codimension one hyperplane in Y_H , and a D' neighborhood of F_0 determines a (Euclidean) cylinder $C \cong \mathbb{D}^2 \times \mathbb{R}^{k-2}$ inside the upper half-plane (C is a Euclidean cylinder since in the upper half-plane hyperbolic balls coincide with Euclidean balls with different centers). This cylinder is at some positive distance above the plane $\mathbb{R}^{k-1} \times \{0\}$. The collection of boundary components B_1, \dots, B_k correspond to boundaries of some horoballs, which manifest as finitely many (Euclidean) balls O_1, \dots, O_k tangent to the plane $\mathbb{R}^{k-1} \times \{0\}$.

A horoball based at some point $(v, 0) \in \mathbb{R}^{k-1} \times \mathbb{R}$ and with Euclidean radius r determines a vertical cylinder $\mathbb{D}^{k-1} \times (0, \infty)$ in the upper half plane, consisting of all points at Euclidean distance $\leq r$ from the line $\{v\} \times (0, \infty)$. Choose a horoball O'' with the property that the corresponding vertical cylinder $C' \cong \mathbb{D}^{k-1} \times \mathbb{R}$ does not intersect the cylinder $C = N_{D'}(F_0)$, and does not intersect any of the $\bar{O}_i = N_D(O_i)$ (the hyperbolic D -neighborhood of the O_i). Let $(v, 0) \in \mathbb{R}^{k-1} \times \{0\}$ be the basepoint for the horoball O'' . Then the intersection of the line $\{(v, t) : 0 < t \leq H\}$ with $\widehat{\mathbb{H}}^k$ provides a path α from the boundary component B to the boundary component $B'' = \partial O''$ which does not intersect $N_{D'}(F_0) \cup N_D(B_1) \cup \dots \cup N_D(B_k)$. If B'' coincides with B' , we are done. Otherwise, we have two distinct boundary components B' and B'' of $\widehat{\mathbb{H}}^k$, which we want to connect in the complement of the D -neighborhood of finitely many other boundary components B_1, \dots, B_k . Choose h as in case (a) small enough so that it additionally satisfies $Y_h \cap C = \emptyset$. Then as in case (a), we know that the corresponding \bar{Y}_h is a manifold with boundary, so we can connect B' to B'' via a path γ in the interior of \bar{Y}_h . This path lies in the complement of $N_D(B_1) \cup \dots \cup N_D(B_k)$, and since $Y_h \cap C = \emptyset$, also lies in the complement of $N_{D'}(F_0)$. Finally, we can choose a path β on the horoball B'' connecting the two endpoints of α and γ . Note that $\beta \subset C'$, while the cylinder C' has the property that $C' \cap (N_{D'}(F_0) \cup N_D(B_1) \cup \dots \cup N_D(B_k)) = \emptyset$. As such, the concatenation $\alpha * \beta * \gamma$ provides a curve joining B to B' in the complement of the set $N_{D'}(F_0) \cup N_D(B_1) \cup \dots \cup N_D(B_k)$, as desired. This completes the argument in case (b), and concludes the “only if” part of the proof.

(\Leftarrow) We are left to show the “if” part. Observe that, since $H \cong \mathbb{Z}^{n-1}$, we have that H is a subgroup of a vertex group (Lemma 13.2). Moreover, since there exists $K \in F(M)$ such that $H \cap K$ is isomorphic to \mathbb{Z}^{n-2} , the piece corresponding to the vertex fixed by H is necessarily a surface piece. Therefore, in order to conclude it suffices to show the following. Suppose that H corresponds (via the Flat Torus

Theorem) to a subspace V of a chamber C_s covering a surface piece. Also, suppose that V is not contained in a wall adjacent to a non-surface piece and let $F \subseteq V$ be a set of the form $\{p\} \times \mathbb{R}^{n-2} \subseteq C_s$. Then for each flat V' (corresponding to some $H' \in F(M)$ with $H' \neq H$), we can identify a $D \geq 0$ and a finite collection of flats V_1, \dots, V_k (corresponding to $H_1, \dots, H_k \in F(M) \setminus \{H', H\}$) with the property that any path β connecting V' to V intersects $N_D(F) \cup N_D(V_1) \cup \dots \cup N_D(V_k)$. We emphasize that the choice of D and the finite family V_i depends on the given V' , but is independent of the path β .

We first argue that, without loss of generality, it is enough to consider the case where V and V' lie in the same surface piece. Indeed, the case where V, V' cannot be chosen in the same surface piece corresponds to the situation where there is no vertex in the Bass-Serre tree T which is fixed by both subgroups $H, H' \in F(M)$. In this case it is easily seen that a wall V_1 in \tilde{M} exists with the property that every path joining V to V' intersects V_1 . The single flat V_1 (with $D = 0$) gives the desired finite collection of flats.

So now we may assume that V, V' lie in the same surface piece C_s . Since these are codimension one flats, they must be of the form $V = \gamma \times \mathbb{R}^{n-2}, V' = \gamma' \times \mathbb{R}^{n-2}$, where γ, γ' are a pair of geodesics in the base $\widehat{\mathbb{H}}^2$ (a copy of the hyperbolic plane \mathbb{H}^2 with a suitable equivariant family of horoballs removed). Moreover, the geodesics γ, γ' project to closed geodesics in the surface with boundary Σ . This implies that the geodesics γ, γ' determine four distinct ends of the space $\widehat{\mathbb{H}}^2$. Now recall that we have a subset $F = \{p\} \times \mathbb{R}^{n-2} \subset V$, where $p \in \gamma$. From the definitions of an end, one can find a suitably large $D > 0$ with the property that the four ends of the geodesics γ, γ' each exit into a different connected component of $\widehat{\mathbb{H}}^2 \setminus N_D(p)$. The metric ball $N_D(p)$ intersects finitely many boundary components of $\widehat{\mathbb{H}}^2$, call them $\gamma_1, \dots, \gamma_k$.

We now claim that the chosen $D > 0$, and finite collection of flats $V_i := \gamma_i \times \mathbb{R}^{n-2}$, has the desired properties. That is to say, any curve β in \tilde{M} joining a point in V to a point in V' must intersect the subset $N_D(F) \cup V_1 \cup \dots \cup V_k$. Arguing by contradiction, let us assume that we can find such a path β in \tilde{M} avoiding the subset $N_D(F) \cup V_1 \cup \dots \cup V_k$. Our first step is to modify the path β so that it lies entirely within the chamber C_s . To do this, we note that the path β can be expressed as a concatenation $\beta_1 * \eta_1 * \beta_2 * \eta_2 * \dots * \eta_r * \beta_r$, where each β_j is a path contained entirely inside the chamber C_s , and each η_j is a path with the property that $\eta_j \cap C_s \subset \partial C_s$. From the tree of space structure on \tilde{M} , we also have that each $\eta_j \cap C_s$ is contained in a *single connected component* W_j of ∂C_s , which by hypothesis must be distinct from the finitely many boundary components V_i , i.e. $W_j \neq V_i$ for all i, j . But the V_i are all the boundary components of C_s which intersect $N_D(F)$, so we conclude that

$$W_j \cap (N_D(F) \cup V_1 \cup \dots \cup V_k) = \emptyset$$

Since each W_j is connected, we can replace each curve η_j by a corresponding curve η'_j , contained in W_j (and hence $\subset C_s$), and having the same endpoints. This yields a new curve β' given by the concatenation $\beta_1 * \eta'_1 * \beta_2 * \eta'_2 * \dots * \eta'_r * \beta_r$, which is now entirely contained inside C_s , joins a point in V to a point in V' , but avoids the set $N_D(F) \cup V_1 \cup \dots \cup V_k$.

Finally, consider the projection $\bar{\beta}$ of the curve β' onto the base $\widehat{\mathbb{H}}^2$. Since all the sets in C_s we were considering were given as products, we easily see that the

curve $\bar{\beta}$ joins a point on the geodesic γ to a point on the geodesic γ' , which avoids the set $N_D(p)$. But recall that D was chosen large enough so that the connected components of $\gamma \setminus N_D(p)$ and $\gamma' \setminus N_D(p)$ each lie in a distinct connected component of $\widehat{\mathbb{H}}^2 \setminus N_D(p)$. This gives us the desired contradiction, and concludes the proof of the Theorem. \square

13.3. Rigidity results

In the last section, we gave a characterization of fundamental groups of internal walls and walls adjacent to a non-surface piece which is easily seen to be invariant under isomorphisms. We can now prove the extended version of Theorem 0.3

THEOREM 13.4 (Isomorphisms preserve pieces). *Let M_1, M_2 be a pair of extended graph manifolds and let $\Gamma_i = \pi_1(M_i)$ be their respective fundamental groups. Let $\Lambda_1 \leq \Gamma_1$ be a subgroup conjugate to the fundamental group of the piece V in M_1 , and $\phi : \Gamma_1 \rightarrow \Gamma_2$ be an isomorphism. Then $\phi(\Lambda_1)$ is conjugate to the fundamental group $\Lambda_2 \leq \Gamma_2$ of a piece in M_2 .*

PROOF. Fundamental groups of internal walls and walls adjacent to a non-surface piece are preserved, by the characterization given in the previous lemma. So, we can proceed as in Proposition 4.3 and Section 4.3 to show that fundamental groups of non-surface pieces are quasi-preserved and hence preserved.

In order to show that fundamental groups of surface pieces are preserved, notice that one such fundamental group is the union of all $H \in F(M)$ containing a fixed subgroup isomorphic to \mathbb{Z}^{n-2} . \square

This provides all the necessary ingredients to establish smooth rigidity.

THEOREM 13.5 (Smooth rigidity for extended graph manifolds). *Let M, M' be extended graph manifolds, and let $\varphi : \pi_1(M) \rightarrow \pi_1(M')$ be a group isomorphism. Suppose that the boundaries of M, M' do not intersect any surface piece. Then φ is induced by a diffeomorphism $\psi : M \rightarrow M'$.*

Notice that the additional hypothesis is necessary, as the result does not hold if M, M' consist of a single surface piece.

PROOF. The argument in Chapter 5 carries over, except that we cannot use Mostow rigidity to deal with surface pieces. However, in the hypotheses of the theorem, we end up dealing with outer automorphisms of surfaces which preserve the conjugacy classes of the fundamental groups of the boundary components. It is well-known that such automorphisms are indeed induced by diffeomorphisms (see e.g. [FaMa, Theorem 8.8]). \square

Let us also sketch the proof of quasi-isometric rigidity. First, here is the statement of quasi-isometry rigidity for surface pieces.

THEOREM 13.6 (QI-rigidity of surface pieces). *Let G be a finitely generated free group, and let Γ be a finitely generated group quasi-isometric to $G \times \mathbb{Z}^d$, $d \geq 0$. Then there exists a finite-index subgroup Γ' of Γ isomorphic to $G' \times \mathbb{Z}^d$, where G' is a free group.*

PROOF. Notice that any asymptotic cone of $G \times \mathbb{Z}^d$ is biLipschitz equivalent to $T \times \mathbb{R}^d$, where T is a tree. By [KaLe1, Lemma 2.14], the image of any biLipschitz embedding of \mathbb{R}^{d+1} into any asymptotic cone of $G \times \mathbb{Z}^d$ has the form $\gamma \times \mathbb{R}^d$, where

γ is a bi-infinite geodesic in T . The cited lemma is actually stated in the case $d = 1$, but the proof applies to this more general case as well (as remarked in [BJN], where the general case is needed). As in Corollary 9.14, one can show that the image of any quasi-isometric embedding of \mathbb{Z}^d into $G \times \mathbb{Z}^d$ is at finite Hausdorff distance from a subset of the form $\gamma \times \mathbb{Z}^d$, where γ is a geodesic in G . Using this ingredient, one can proceed as in Chapter 10 using the well-known fact that a group quasi-isometric to a free group is virtually free (this is essentially due to Stallings [St] and Dunwoody [Du]) instead of the results by Schwartz.

This gives us a finite-index subgroup Γ'' of Γ and a free group Δ such that the following short exact sequence holds:

$$1 \longrightarrow A \xrightarrow{j} \Gamma'' \longrightarrow \Delta \longrightarrow 1,$$

where A contains \mathbb{Z}^d as a finite index subgroup and $j(A)$ is contained in the center of Γ'' . Therefore, as Δ is free, we have that Γ'' is isomorphic to $\Delta \times A$, so Γ'' contains a finite index subgroup Γ' isomorphic to $\mathbb{Z}^d \times \Delta$, and we are done. \square

We have to show that quasi-isometries of the universal cover of irreducible extended graph manifolds quasi-preserve chambers. We can assume that we are dealing with manifolds of dimension at least 4, as in the case of dimension 3 we would be dealing with classical graph manifolds, and this fact is already known in this case (see [KaLe3]). The advantage is that under this assumption we do not have adjacent surface pieces (otherwise, fiber subgroups of adjacent pieces should intersect non-trivially, contradicting irreducibility).

THEOREM 13.7 (QI's preserve pieces of irreducible graph manifolds, extended case). *Let M_1, M_2 be a pair of irreducible extended graph manifolds, and $\Gamma_i = \pi_1(M_i)$ their respective fundamental groups. Let $\Lambda_1 \leq \Gamma_1$ be a subgroup conjugate to the fundamental group of a piece in M_1 , and $\varphi : \Gamma_1 \rightarrow \Gamma_2$ be a quasi-isometry. Then, the set $\varphi(\Lambda_1)$ is within finite Hausdorff distance from a conjugate of $\Lambda_2 \leq \Gamma_2$, where Λ_2 is the fundamental group of a piece in M_2 .*

SKETCH OF PROOF. First, the proof of Theorem 7.11 applies to our case as well, so we have that chambers are quasi-isometrically embedded.

Define a *barrier* to be a subset of \widetilde{M}_i which is either a wall separating non-surface pieces or a chamber covering a surface piece. As in Lemma 9.12, we can show that distinct ω -barriers have the property that all paths connecting them pass through a subset homeomorphic to \mathbb{R}^{n-3} . In particular, as in Proposition 9.13, we have that the image of each topological embedding of \mathbb{R}^{n-1} into an asymptotic cone of \widetilde{M}_i is *contained* in an ω -barrier. Using this, one can show that the image of each quasi-isometric embedding of \mathbb{R}^{n-1} into \widetilde{M}_i (henceforth, quasi-flat) is contained in a finite radius neighborhood of a barrier. Notice that each barriers covering a surface piece can be characterized, up to finite Hausdorff distance, as the union of all quasi-flats (for a given quasi-isometry constant) whose coarse intersection is a specified subset quasi-isometric to \mathbb{R}^{n-2} . Combining this with the fact that walls are quasi-preserved, we get that barriers are quasi-preserved. So, chambers covering surface pieces are quasi-preserved, and we are left to show the same result for the other chambers as well. In view of the proof of Proposition 4.3, we only need to show that walls adjacent to chambers covering non-surface pieces are preserved. This follows from the obvious metric analogue of Theorem 13.3. \square

The quasi-preservations of chambers was the key fact to reduce the study of a group quasi-isometric to the fundamental group of a graph manifold to that of groups quasi-isometric to fundamental groups of pieces. In fact, the same argument we used in Section 11 gives the following.

THEOREM 13.8. *Let M be an irreducible graph n -manifold obtained by gluing the pieces $V_i = \overline{N}_i \times T^{d_i}$, $i = 1, \dots, k$. Let Γ be a group quasi-isometric to $\pi_1(M)$. Then either Γ itself or a subgroup of Γ of index two is isomorphic to the fundamental group of a graph of groups satisfying the following conditions:*

- every edge group contains \mathbb{Z}^{n-1} as a subgroup of finite index;
- every vertex group is either of the form described by Theorem 0.13 or it is, up to finite index, the direct product of a free group and \mathbb{Z}^{n-1} .

Directions for future research

Our purpose in this monograph was to initiate the study of the class of high-dimensional graph manifolds. In this final chapter, we collate various problems that came up naturally in the course of this work, and could serve as directions for future research.

14.1. Further algebraic properties

In Chapter 6, we established various algebraic properties of the fundamental groups of high dimensional graph manifolds. Most of the results followed fairly easily from the structure of such groups, expressed as a graph of groups. In Chapter 8, we obtained some additional results in the case where the graph manifold was assumed to be irreducible. Most of these were consequences of the fact that, in the irreducible case, the associated action on the Bass-Serre tree was acylindrical. In contrast, there are a number of interesting properties of groups whose behavior under amalgamations is less predictable. It would be interesting to see which of these properties hold for the class of graph manifold groups. For concreteness, we identify some properties which we think would be of most interest:

PROBLEM. Are fundamental groups of high dimensional graph manifolds Hopfian? Are they residually finite? Are they linear? What if one additionally assumes the graph manifold is irreducible?

A slightly different flavor of problems come from the algorithmic viewpoint. We showed that the word problem is solvable for the $\pi_1(M)$ of irreducible graph manifolds. Some other algorithmic problems one can consider include:

PROBLEM. Is the conjugacy problem solvable for fundamental groups of high dimensional graph manifolds? Is the isomorphism problem solvable within the class of graph manifold groups?

Finally, one can also ask for a better understanding of the outer automorphism group $\text{Out}(\pi_1(M))$, and of how it relates to the topology of M . For instance:

PROBLEM. Is the group $\text{Out}(\pi_1(M))$ always infinite? What can be said about the structure of $\text{Out}(\pi_1(M))$?

PROBLEM. If we have a finite subgroup in $\text{Out}(\pi_1(M))$, can we lift it back to a finite subgroup of $\text{Diff}(M)$?

This last problem is an analogue of the classic Nielsen realization problem. Note that, by Theorem 0.5, the natural map $\text{Diff}(M) \rightarrow \text{Out}(\pi_1(M))$ is surjective. So we can always lift back individual elements from $\text{Out}(\pi_1(M))$ to $\text{Diff}(M)$, and the problem asks whether we can choose the lifts in a compatible manner.

14.2. Studying quasi-isometries

One of our main results, Theorem 0.14, gives us some structure theory for groups which are quasi-isometric to the fundamental group of an irreducible graph manifold. Specializing to the class of graph manifold groups, this result gives us a necessary condition for deciding whether two such groups $\pi_1(M_1)$ and $\pi_1(M_2)$ are quasi-isometric to each other: loosely speaking, the two graph manifolds M_i must essentially be built up from the same collection of pieces (up to commensurability), with the same patterns of gluings. The only distinguishing feature between M_1 and M_2 would then be in the actual gluing maps used to attach pieces together. This brings us to the interesting:

PROBLEM. *To what extent do the gluing maps influence the quasi-isometry type of the resulting graph manifold group? More concretely, take pieces V_1 and V_2 each having exactly one boundary component, and let M_1, M_2 be a pair of irreducible graph manifolds obtained by gluing V_1 with V_2 . Must the the fundamental groups of M_1 and M_2 be quasi-isometric?*

In order to prove that the answer is positive, one could try to follow the strategy described in [BeNe], as follows:

- (1) Define a *flip manifold* as a graph manifold whose gluing maps are such that fibers are glued to parallel copies of the traces at the toric boundaries of the adjacent base (this definition generalizes the one given in [KaLe4]).
- (2) Observe that since V_1 and V_2 can be glued to provide irreducible graph manifolds, they can also be glued to obtain a flip manifold M . Note however that such a manifold is not uniquely determined by V_1 and V_2 .
- (3) Prove that the universal covering of M_i , $i = 1, 2$, is quasi-isometric to the universal covering of M .

The analogue of Step (3) for pieces with 2-dimensional bases is proved in Section 2 of [KaLe4]. However, the argument given there does not apply in our case, since our bases are not negatively curved.

In Theorem 11.8, we argued that a *labelled* version of the Bass-Serre tree associated to an irreducible graph manifold (with each vertex labelled by the commensurability class of the hyperbolic factor in the corresponding vertex group) provides a quasi-isometric invariant. However, it is showed in Remark 11.10 that this is *not* a complete invariant, i.e. that there exist a pair of irreducible graph manifolds with the same invariant, but which are nevertheless not quasi-isometric. We can ask:

PROBLEM. *Can one devise a more sophisticated labeling in order to get a complete quasi-isometric invariant?*

It would be interesting to see how the quasi-isometry classes behave with respect to curvature conditions. For instance, we could ask:

PROBLEM. *Is there a pair of irreducible graph manifolds with quasi-isometric fundamental groups, with the property that one of them supports a locally CAT(0) metric, but the other one cannot support any locally CAT(0) metric?*

Note that if the quasi-isometry class ends up being independent of the gluing maps used (among the ones giving irreducible graph manifolds), then by varying the gluing maps, one can give an affirmative answer to this last question.

Now all the quasi-isometry results we have are for the class of irreducible graph manifolds. The key result we use is that, for this class of graph manifolds, all the walls are *undistorted* in the universal cover (see Chapter 7, particularly the discussion in Remark 7.14). This in turn can be used to show that quasi-isometries must send walls to walls (up to finite distance), and hence chambers to chambers (see Chapter 9). Trying to generalize these, we can formulate the following question, which was suggested to us by C. Drutu and P. Papasoglu:

PROBLEM. *For a graph manifold M , assume that a wall W in the universal cover \tilde{M} is not too distorted (say, polynomially distorted). What additional hypotheses are sufficient to ensure that quasi-isometries send walls to (bounded distance from) walls? And how can we choose gluings in order to ensure these hypotheses are satisfied?*

For example, one possibility is to assume that all fibers have dimension which is small relative to the degree of polynomial growth. It seems like this constraint might be enough to show that walls are rigid under quasi-isometries. Finally, we have the most general (and consequently, the most difficult):

PROBLEM. *Develop methods to analyze quasi-isometries of general graph manifolds (i.e. without the assumption of irreducibility).*

Notice that in the proof of Theorem 0.14 we studied each vertex stabilizer separately. It might be possible to obtain additional information by studying the interaction between vertex stabilizers of adjacent vertices.

PROBLEM. *Is it possible, under additional hypotheses, to obtain a better description of the vertex stabilizers?*

A possible strategy to achieve this is to use the fact that walls admit “foliations” which are coarsely invariant under quasi-isometries, namely those given by fibers of the adjacent chambers. In order to obtain additional information out of this, one probably has to assume that the dimension of the fibers is half that of the walls.

14.3. Non-positive curvature and differential geometry

We have already given three different constructions of high dimensional graph manifolds which cannot support a locally CAT(0) metric (see Section 2.5 and Chapter 12), and hence no Riemannian metric of non-positive sectional curvature. It would be interesting to identify precise conditions for such metrics to exist:

PROBLEM. *Find necessary and sufficient conditions for a graph manifold M to*

- (i) support a Riemannian metric of non-positive sectional curvature, or*
- (ii) support a locally CAT(0)-metric.*

It is not even clear whether or not items (i) and (ii) above are really distinct:

PROBLEM. *Assume the high dimensional graph manifold M supports a locally CAT(0) metric. Does it follow that M supports a Riemannian metric of non-positive sectional curvature?*

Note that, for the classical 3-dimensional graph manifolds, Buyalo and Svetlov [BuSv] have a complete criterion for deciding whether or not such a manifold supports a non-positively curved Riemannian metric (see also [Le]). Some partial results in dimension = 4 appear in [BuKo].

Concerning the second problem, in the 3-dimensional setting, there is no difference between Riemannian and metric non-positive curvature (see for instance [DJL, Section 2]). However, in all dimensions ≥ 4 , there exist manifolds supporting locally CAT(0) metrics which do *not* support Riemannian metrics of non-positive curvature (see the discussion in [DJL, Section 3]). For the class of graph manifolds, the situation is relatively tame, and one might expect the two classes to coincide.

Next, we discuss a question about ordinary hyperbolic manifolds. One can ask whether examples exist satisfying a strong form of the cohomological condition appearing in Proposition 12.4. More precisely:

PROBLEM. *Can one find, in each dimension $n \geq 4$, an example of a truncated finite volume hyperbolic n -manifold N , with all boundary components consisting of tori, such that at least one boundary component T has the property that the map $i_* : H_1(T) \rightarrow H_1(N)$ induced by inclusion has a non-trivial kernel?*

Note that such examples clearly exist in dimensions $= 2, 3$. Moreover, if one could construct a finite volume hyperbolic n -manifold having a single cusp with toral cross section, then Proposition 12.4 could be used to show that the corresponding $\ker(i_*)$ is non-trivial. The problem of constructing hyperbolic manifolds with a single cusp is, however, still open.

Our next question comes from a differential geometric direction. Intuitively, one can think of high dimensional graph manifolds as being “mostly” non-positively curved: the difficulties in putting a global metric of non-positive curvature is concentrated in the vicinity of the gluing tori, which are a collection of smooth, pairwise disjoint, codimension one submanifolds. Gromov has formulated the notion of *almost non-positively curved manifolds*: these are manifolds with the property that for each $\epsilon > 0$, one can find a Riemannian metric with the property that the diameters d and maximal sectional curvature K satisfy the inequality $K \cdot d^2 \leq \epsilon$ (see [Gr3]). It would be interesting to study graph manifolds from this viewpoint. In particular:

PROBLEM. *Are graph manifolds almost non-positively curved?*

We note that the class of almost non-positively curved manifolds is very mysterious. The only known examples of manifold which are known to *not* be almost non-positively curved are the sphere S^2 and the projective plane $\mathbb{R}P^2$ (by Gauss-Bonnet). Aside from manifolds supporting non-positive curvature, the only additional known examples of almost non-positively curved manifolds occur in dimension $= 3$ (all 3-manifolds are non-positively curved, see Bavard [Ba]) and in dimension $= 4$ (a family of examples was constructed by Galaz-Garcia [G-G]).

Keeping on the theme of differential geometry, we recall that the minimal volume of a smooth manifold is defined to be the infimum of the volume functional, over the space of all Riemannian metrics whose curvature is bounded between -1 and 1 . Gromov [Gr4] showed that manifolds with positive simplicial volume have positive minimal volume and have positive minimal entropy. In view of our Proposition 6.8, one can ask the following:

PROBLEM. *Let M be a graph manifold with at least one purely hyperbolic piece (i.e. a piece with trivial fiber). Can one compute the minimal volume of M ? Does*

it equal the sum of the hyperbolic volumes of the purely hyperbolic pieces? Does the choice of gluing maps between tori affect this invariant? If there are some pieces with non-trivial fiber, can the minimal volume ever be attained by an actual metric on M ?

Similarly, minimal entropy is defined to be the infimum of the topological entropy of the geodesic flow, over the space of all Riemannian metrics whose volume is equal to one. Gromov [Gr4] also showed that positive simplicial volume implies positive minimal entropy. One could formulate the same types of questions concerning the minimal entropy.

Bibliography

- [A-K] Y. Algom-Kfir, *Strongly contracting geodesics in outer space*, preprint available on the arxiv:math/0812.1555.
- [AlBr] J. M. Alonso & M. R. Bridson, *Semihyperbolic groups* Proc. Lond. Math. Soc. **70** (1995), 56-114.
- [ArFa] C. S. Aravinda & F. T. Farrell, *Twisted doubles and nonpositive curvature*, Bull. Lond. Math. Soc. **41** (2009), 1053-1059.
- [AM] G. Arzhantseva & A. Minasyan, *Relatively hyperbolic groups are C^* -simple*, J. Funct. Anal., **243** (2007), 345-351.
- [AMO] G. Arzhantseva, A. Minasyan & D. Osin, *The SQ-universality and residual properties of relatively hyperbolic groups*, J. Algebra, **315** (2007), 165-177.
- [AO] G. Arzhantseva & D. Osin, *Solvable groups with polynomial Dehn functions*, Trans. Amer. Math. Soc. **354** (2002), 3329-3348.
- [BaLu] A. Bartels & W. Lück, *Isomorphism conjecture for homotopy K-theory and groups acting on trees*, J. Pure Appl. Algebra **205** (2006), 660-696.
- [Ba] H. Bass, *Covering theory for graphs of groups*, J. Pure Appl. Algebra **89** (1993), 3-47.
- [BHS] H. Bass, A. Heller, & R. G. Swan, *The Whitehead group of a polynomial extension*, Inst. Hautes Études Sci. Publ. Math. **22** (1964), 61-79.
- [Ba] C. Bavard, *Courbure presque négative en dimension 3*, Compos. Math. **63** (1987), 223-236.
- [BdlHV] B. Bekka, P. de la Harpe & A. Valette, *Kazhdan's property (T)*. New Mathematical Monographs, **11**. Cambridge University Press, Cambridge, 2008.
- [BDM] J. A. Behrstock, C. Druţu & L. Mosher *Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity*, Math. Ann. **344** (2009), 543-595.
- [BDS] J. A. Behrstock, C. Druţu & M. Sapir, *Median structures on asymptotic cones and homomorphisms into mapping class groups*, to appear in Proc. Lond. Math. Soc.
- [BJN] J. A. Behrstock, T. Januszkiewicz & W. D. Neumann, *Quasi-isometric classification of some high dimensional right-angled Artin groups*. Groups Geom. Dyn. **4** (2010), no. 4, 681-692.
- [BeNe] J. A. Behrstock & W. D. Neumann, *Quasi-isometric classification of graph manifold groups* Duke Math. J. **141** (2008), 217-240.
- [BePe] R. Benedetti, C. Petronio, *Lectures on hyperbolic geometry*. Universitext. Springer-Verlag, Berlin, 1992.
- [Be] A. A. Bernasconi, *On HNN-extensions and the complexity of the word problem for one relator groups*, Ph.D. Thesis, University of Utah, June 1994. Available online at <http://www.math.utah.edu/~sg/Papers/bernasconi-thesis.pdf>
- [Br] N. Brady, *Branched coverings of cubical complexes and subgroups of hyperbolic groups*, J. Lond. Math. Soc. **60** (1999), 461-480.
- [Br-dlH] M. Bridson & P. de la Harpe, *Mapping class groups and outer automorphism groups of free groups are C^* -simple*, J. Funct. Anal. **212** (2004), 195-205.
- [BGHM] M. Bridson, D. Groves, J. A. Hillman, & G. J. Martin, *Cofinitely Hopfian groups, open mappings and knot complements*, Groups Geom. Dyn. **4** (2010), 693-707.
- [BrHa] M. R. Bridson & A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren der Mathematischen Wissenschaften **319**. Springer-Verlag, Berlin, 1999.
- [Bu-dlH] M. Bucher & P. de la Harpe, *Mapping class groups and outer automorphism groups of free groups are C^* -simple*, Math. Notes **67** (2000), 686-689.
- [BuTa] J. Burillo & J. Taback, *Equivalence of geometric and combinatorial Dehn functions*, New York J. Math. **8** (2002), 169-179.

- [BuKo] S. V. Buyalo & V. L. Kobelskiĭ, *Generalized graph-manifolds of nonpositive curvature*, St. Petersburg Math. J. **11** (2000), 251-268.
- [BuSv] S. V. Buyalo & P. Svetlov, *Topological and geometric properties of graph-manifolds*, St. Petersburg Math. J. **16** (2005), 297-340.
- [Ca] S. Cappell, *A splitting theorem for manifolds*, Invent. Math. **33** (1976), pp. 69-170.
- [CCJJV] P.-A. Cherix, M. Cowling, P. Jolissaint, P. Julg, & A. Valette, *Groups with the Haagerup property*. Progress in Mathematics, **197**, Birkhäuser Verlag, Basel, 2001. viii+126 pp.
- [CoPr] F. X. Connolly & S. Prassidis, *On the exponent of the cokernel of the forget-control map on K_0 -groups*, Fund. Math. **172** (2003), pp. 201-216.
- [Da] F. Dahmani, *Combination of convergence groups*, Geom. Topol. **7** (2003), 933-963.
- [DJL] M. Davis, T. Januszkiewicz, & J.-F. Lafont, *4-dimensional locally $CAT(0)$ -manifolds with no Riemannian smoothings*, preprint available on the arXiv:1002.4235
- [dlH] P. de la Harpe, *On simplicity of reduced C^* -algebras of groups*, Bull. Lond. Math. Soc. **39** (2007), 1-26.
- [De] T. Delzant, *Sur l'accessibilité acylindrique des groupes de présentation finie*, Ann. Inst. Fourier **49** (1999), 1215-1224.
- [Do] A. Dold, *A simple proof of the Jordan-Alexander complement theorem*, Amer. Math. Monthly **100** (1993), 856-857.
- [DMS] C. Druţu, S. Mozes & M. Sapir, *Divergence in lattices in semisimple Lie groups and graphs of groups*, Trans. Amer. Math. Soc. **362** (2010), 2451-2505.
- [DrSa] C. Druţu & M. Sapir, *Tree-graded spaces and asymptotic cones of groups. With an appendix by D. Osin and M. Sapir*, Topology **44** (2005), 959-1058.
- [Du] M. J. Dunwoody, *The accessibility of finitely presented groups*, Invent. Math. **81** (1985), no. 3, 449-457.
- [EMO] A. Eskin, S. Mozes, & H. Oh, *On Uniform Exponential Growth for Linear Groups*, Int. Math. Res. Not. (2002), no. 31, 1675-1683
- [Fa1] B. Farb, *The Extrinsic geometry of subgroups and the generalized word problem*, Proc. Lond. Math. Soc. **68** (1994), 577-593.
- [Fa2] B. Farb, *Relatively hyperbolic groups*, Geom. Funct. Anal. **8** (1998), 810-840.
- [FaMa] B. Farb & D. Margalit, *A primer on mapping class groups*, available at <http://www.math.uchicago.edu/~margalit/mcg/mcgv50.pdf>
- [Fa] F. T. Farrell, *Surgical methods in rigidity*. Springer-Verlag, Berlin, 1996. iv + 98 pp.
- [FaJo1] F. T. Farrell & L. E. Jones, *Topological rigidity for compact non-positively curved manifolds*, in "Differential geometry: Riemannian geometry (Los Angeles, CA, 1990)", Proc. Sympos. Pure Math. **54** (1993), 229-274.
- [FaJo2] F. T. Farrell & L. E. Jones, *Rigidity for aspherical manifolds with $\pi_1 \subset GL_m(\mathbf{R})$* , Asian J. Math. **2** (1998), 215-262.
- [G-G] F. Galaz-Garcia, *Examples of 4-manifolds with almost nonpositive curvature*, Differential Geom. Appl. **26** (2008), 697-703.
- [Ge] S. M. Gersten, *Isoperimetric and isodiametric functions*, in "Geometric Group Theory I", ed. by G. Niblo and M. Roller. LMS Lecture notes **181**. Cambridge Univ. Press, 1993.
- [Ge2] S. M. Gersten, *Quadratic divergence of geodesics in $CAT(0)$ spaces*, Geom. Funct. Anal. **4** (1994), 37-51.
- [GHH] O. Goodman, D. Heard, & C. Hodgson, *Commensurators of cusped hyperbolic manifolds*, Experiment. Math. **17** (2008), 283-306.
- [Gr1] M. Gromov, *Groups of polynomial growth and expanding maps*, Inst. Hautes Études Sci. Publ. Math. **53** (1981), 53-73.
- [Gr2] M. Gromov, *Hyperbolic groups*, in "Essays in group theory", 75-263, Math. Sci. Res. Inst. Publ. **8** (1987), Springer, New York.
- [Gr3] M. Gromov, *Almost flat manifolds*, J. Differential Geom. **13** (1978) 231-241.
- [Gr4] M. Gromov, *Volume and bounded cohomology*, Inst. Hautes Études Sci. Publ. Math. **56** (1982), 5-99.
- [GeSh] S. M. Gersten & H. Short, *Rational subgroups of biautomatic groups*, Ann. Math. **134** (1991), 125-158.
- [Ha] M. Hall, *Finiteness conditions for soluble groups*, Proc. Lond. Math. Soc. **4** (1954), 419-436.

- [HK] N. Higson & G. Kasparov, *E-theory and KK-theory for groups which act properly and isometrically on Hilbert space*, Invent. Math. **144** (2001), 23-74.
- [HsWa] W. C. Hsiang & C. T. C. Wall, *On homotopy tori II*, Bull. Lond. Math. Soc. **1** (1969), 341-342.
- [Ka] I. Kapovich, *The combination theorem and quasiconvexity*, Internat. J. Algebra Comput. **11** (2001), 185-216.
- [KaLe1] M. Kapovich & B. Leeb, *On asymptotic cones and quasi-isometry classes of fundamental groups of 3-manifolds*, Geom. Funct. Anal. **5** (1995), no. 3, 582-603.
- [KaLe2] M. Kapovich & B. Leeb, *Actions of discrete groups on nonpositively curved spaces*, Math. Ann. **306** (1996), 341-352.
- [KaLe3] M. Kapovich & B. Leeb, *Quasi-isometries preserve the geometric decomposition of Haken manifolds*, Invent. Math. **128** (1997), 393-416.
- [KaLe4] M. Kapovich & B. Leeb, *3-manifold groups and non-positive curvature*, Geom. Funct. Anal. **8** (1998), 841-852.
- [KILe] B. Kleiner & B. Leeb, *Groups quasi-isometric to symmetric spaces*, Comm. Anal. Geom. **9** (2001), 239-260.
- [Ko] M. Koubi, *Croissance uniforme dans les groupes hyperboliques*, Ann. Inst. Fourier **48** (1998), 1441-1453.
- [Ku] T. Kuessner, *Multicomplexes, bounded cohomology and additivity of the simplicial volume*, preprint available on the arxiv:math/0109057v2.
- [La] J.-F. Lafont, *A boundary version of Cartan-Hadamard and applications to rigidity*, J. Topol. Anal. **1** (2009), 431-459.
- [LoRe] D. D. Long & A. W. Reid, *All flat manifolds are cusps of hyperbolic orbifolds*, Algebr. Geom. Topol. **2** (2002), 285-296.
- [Le] B. Leeb, *3-manifolds with(out) metrics of nonpositive curvature*, Invent. Math. **122** (1995), 277-289.
- [LuR] W. Lück & H. Reich, *The Baum-Connes and the Farrell-Jones conjectures in K- and L-theory*, in "Handbook of K-theory. Vol. 1, 2", pp. 703-842. Springer, Berlin, 2005.
- [LySc] R. C. Lyndon & P. E. Schupp, *Combinatorial group theory*. Reprint of the 1977 edition. Classics in Mathematics. Springer-Verlag, Berlin, 2001. xiv+339 pp.
- [Ma] A. Manning, *Topological entropy for geodesic flows*, Ann. of Math. **110** (1979), 567-573.
- [MRS] D. B. McReynolds, A. Reid, & M. Stover, *Collisions at infinity in hyperbolic manifolds*, in preparation.
- [MNS] C. F. Miller III, W. D. Neumann, & G. Swarup, *Some examples of hyperbolic groups*, in "Geometric group theory down under (Canberra, 1996)", 195-202, de Gruyter, Berlin, 1999.
- [MV] G. Mislin & A. Valette, *Proper group actions and the Baum-Connes conjecture*, Advanced Courses in Mathematics - CRM Barcelona, Birkhäuser Verlag, Basel, 2003. viii+131 pp.
- [MSW1] L. Mosher, M. Sageev, & K. Whyte, *Quasi-actions on trees I. Bounded Valence*, Ann. of Math. (2) **158** (2003), no. 1, 115-164.
- [MSW2] L. Mosher, M. Sageev, & K. Whyte, *Quasi-actions on trees II: Finite depth Bass-Serre trees*, preprint available on the arXiv: math.GR/0405237
- [Ng] T. T. Nguyen Phan, *Smooth (non)rigidity of cusp-decomposable manifolds*, to appear in Comment. Math. Helv.
- [NeRe] W.D. Neumann & L. Reeves, *Central extensions of word hyperbolic groups*, Ann. Math. **145** (1997), 183-192.
- [Ol] A. Yu. Olshanskii, *SQ-universality of hyperbolic groups*, Mat. Sb. **186** (1995) 119-132 (in Russian); Sb. Math. **186** (1995) 1199-1211.
- [On] P. Ontaneda, *The double of a hyperbolic manifold and non-positively curved exotic PL structures*, Trans. Amer. Math. Soc. **355** (2002), 935-965.
- [Os] D. V. Osin, *Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems*, Mem. Amer. Math. Soc. **179** (2006), no. 843.
- [O-O] H. Oyono-Oyono, *Baum-Connes conjecture and group actions on trees*, K-Theory **24** (2001), 115-134.
- [Pa] P. Papasoglu, *Group splittings and asymptotic topology*, J. Reine Angew. Math. **602** (2007), 1-16.

- [Po] R. T. Powers, *Simplicity of the C^* -algebra associated with the free group on two generators*, Duke Math. J. **42** (1975), 151-156.
- [PaVa] I. Pays & A. Valette, *Sous-groupes libres dans les groupes d'automorphismes d'arbres*, Enseign. Math. **37** (1991), 151-174.
- [Ru] K. Ruane, *CAT(0) boundaries of truncated hyperbolic space*, Spring Topology and Dynamical Systems Conference. Topology Proc. **29** (2005), 317-331.
- [Sc] R. E. Schwartz, *The quasi-isometry classification of rank-one lattices*, Inst. Hautes Études Sci. Publ. Math. **82** (1995), 133-168.
- [Se] J.-P. Serre, *Trees*. Springer monographs in Mathematics, Springer-Verlag, Berlin and New York, 1980.
- [Si] A. Sisto, *Projections and relative hyperbolicity*, preprint available on the arXiv:1010.4552.
- [St] J. R. Stallings, *On torsion-free groups with infinitely many ends*, Ann. of Math. (2) **88** (1968), 312-334.
- [Wa1] F. Waldhausen, *On irreducible 3-manifolds which are sufficiently large*, Ann. of Math. **87** (1968), 56-88.
- [Wa2] F. Waldhausen, *Whitehead groups of generalized free products*, in "Algebraic K-theory, II," pp. 155-179. Lecture Notes in Math., Vol. **342**, Springer, Berlin, 1973.
- [Wa3] F. Waldhausen, *Algebraic K-theory of generalized free products. I, II*, Ann. Math. **108** (1978), 135-204.
- [Wa4] F. Waldhausen, *Algebraic K-theory of generalized free products. III, IV*, Ann. Math. **108** (1978), 205-256.
- [Wal1] C. T. C. Wall, *Surgery on compact manifolds*. Amer. Math. Soc., Providence, 1999.
- [Wal2] C. T. C. Wall, *The geometry of abstract groups and their splittings*, Rev. Mat. Complut. **16** (2003), 5-101.
- [Wh] J. H. C. Whitehead, *On the asphericity of regions in a 3-sphere*, Fund. Math. **32** (1939), 149-166.
- [vdDWi] L. van den Dries & J. Wilkie, *Gromov's theorem on groups of polynomial growth and elementary logic*, J. Algebra **89** (1984), 349-374.
- [Zi] R. Zimmer, *Ergodic theory and semi-simple Lie groups*. Birkhauser, Boston, 1984.
- [www] http://www.ms.unimelb.edu.au/~snap/commens/comm_grouped