J. Inst. Math. Jussieu (2022), 1–11 1
doi:10.1017/S1474748022000329

PRESCRIBED VIRTUAL HOMOLOGICAL TORSION OF 3-MANIFOLDS

MICHELLE CHU AND DANIEL GROVES

Department of Mathematics, Statistics, and Computer Science, University of Illinois at Chicago, 322 Science and Engineering Offices (M/C 249), 851 S. Morgan St., Chicago, IL 60607-7045
(michu@uic.edu; groves@math.uic.edu)

(Received 25 June 2021; revised 9 May 2022; accepted 11 May 2022)

Abstract Let M be an irreducible 3-manifold M with empty or toroidal boundary which has at least one hyperbolic piece in its geometric decomposition, and let A be a finite abelian group. Generalizing work of Sun [20] and of Friedl–Herrmann [7], we prove that there exists a finite cover $M' \to M$ so that A is a direct factor in $H_1(M',\mathbb{Z})$.

1. Introduction

In [20], Sun showed that any closed hyperbolic 3-manifold virtually contains any prescribed finite subgroup in homological torsion. Sun used the immersed almost-Fuchsian surfaces of Kahn and Markovic [14] to construct immersed π_1-injective 2-complexes. By using Agol's result that the fundamental groups of closed hyperbolic 3-manifolds are virtually compact special [2] and the implications on virtual retractions to quasi-convex subgroups, Sun finds for any closed hyperbolic 3-manifold a finite cover containing the prescribed finite abelian group as a direct factor in homology [20, Theorem 1.5].

Since the Kahn–Markovic construction requires that the manifolds be closed, Sun's results do not apply to hyperbolic 3-manifolds with cusps. Indeed, Sun asked whether his result applied also to finite-volume hyperbolic 3-manifolds with cusps. In this paper, we extend the results of Sun to a larger class of 3-manifolds which includes all finite-volume hyperbolic 3-manifolds, giving a positive answer to [20, Question 1.8].

Theorem 1.1. *Suppose that M is an irreducible 3-manifold with empty or toroidal boundary which has at least one hyperbolic piece in its geometric decomposition and that A is a finite abelian group. There is a finite cover $N \to M$ so that $H_1(N;\mathbb{Z})$ has a direct factor isomorphic to A.*

Prior to Theorem 1.1, Friedl and Herrmann used [20] and a result of Hadari [10] to show that for any such M and any $k > 0$ there is finite cover $N \to M$ with $|\mathrm{Tor} H_1(N;\mathbb{Z})| > k$ [7, Theorem 1.3]. Independently, Liu showed that any such M admits a finite cover $N' \to M$ with $|\mathrm{Tor} H_1(N';\mathbb{Z})| \neq 0$ [16, Corollary 1.4].

The key case in the proof of Theorem 1.1 is the case of finite-volume hyperbolic 3-manifolds. We follow the strategy of [20] but give an independent proof which simplifies and generalizes Sun's results. We replace Sun's use of the results of Kahn and Markovic [14] with those of Kahn and Wright [15] and replace some arguments of Sun with an elementary argument using coverings of surfaces. We then apply virtual retraction properties of relatively quasi-convex subgroups in relatively hyperbolic groups to deduce both the case where M is finite-volume hyperbolic and also to reduce the general case to this one. Cooper and Futer [6] independently obtained similar results to those of [15] on constructing many closed immersed π_1-injective quasi-Fuchsian surfaces in finite-volume hyperbolic 3-manifolds with cusps. However, our arguments rely on the additional control on the quasi-conformal constants and the holonomies in the Kahn–Wright constructions.

A *hybrid* hyperbolic manifold is constructed either by inbreeding (c.f. [1, 4]) or interbreeding (c.f. [8]) arithmetic hyperbolic manifolds. For $n > 3$ every arithmetic hyperbolic n-manifold N of simplest type contains a totally geodesic arithmetic hyperbolic 3-manifold M (coming from restrictions of the associated quadratic form). By [3, §9], we get the following corollary (some of these cases follow from [20]).

Corollary 1.2. *Suppose that $n > 3$ and N is a finite-volume hyperbolic n-manifold which is either arithmetic of simplest type or a hybrid. Then if A is a finite abelian group, there is a finite cover $N_1 \to N$ so that $H_1(N_1;\mathbb{Z})$ has a direct factor isomorphic to A.*

2. Quasi-isometric embeddings

Our first goal is to prove Theorem 1.1 in the case of a noncompact finite-volume hyperbolic 3-manifold. In Section 6, we deduce the general case from this case. In this section, we record some elementary facts about quasi-isometries and hyperbolic spaces.

Definition 2.1. Let k,λ,κ be constants, and let X,Y be metric spaces. A map $f\colon X \to Y$ is a k–local (λ,κ)-quasi-isometric embedding if for all $x \in X$ the map

$$f|_{B_k(x)}\colon B_k(x) \to Y$$

is a (λ,κ)-quasi-isometric embedding.

The following is essentially [15, Theorem A.20].

Proposition 2.2. *For all δ, for all $\kappa \geq 0$ and all $\lambda \geq 1$, there exist k,λ',κ' so that if Y is a δ–hyperbolic metric space and X is a geodesic metric space, then any k-local (λ,κ)-quasi-isometric embedding $f\colon X \to Y$ is a (λ',κ')-quasi-isometric embedding.*

Proof. Since X and Y are geodesic metric spaces, distances in X and Y are calculated by geodesics. Therefore, we can apply the standard local-to-global result for quasi-geodesics (see, for example, [5, Theorem 3.1.4, p.25]). □

2.1. Half-planes

Let $\theta \in (0,\pi]$. The space P_θ is the subspace of $\mathbb{H}^3$ obtained from gluing two totally geodesic half-planes together along their boundary geodesic, meeting at angle θ. There is a natural

https://doi.org/10.1017/S1474748022000329 Published online by Cambridge University Press

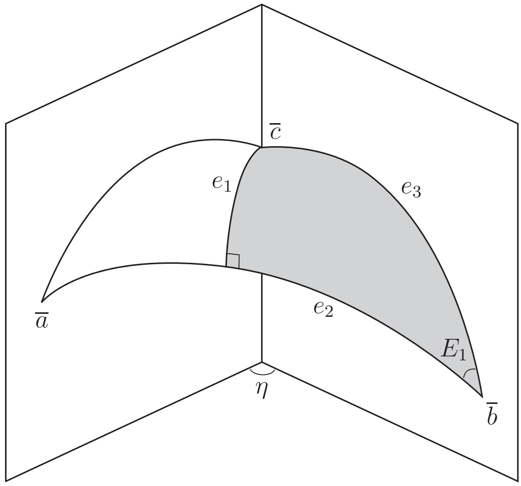

Figure 1. The proof of Lemma 2.3

embedding $p_\theta\colon \mathbb{H}^2 \to \mathbb{H}^3$ given by mapping the positive y-axis to the boundary geodesic of the two half-planes (we consider $\mathbb{H}^2$ in the upper half-space model as a subset of $\mathbb{R}^2$). The image of these boundary geodesics is the *pleating locus* for p_θ.

Lemma 2.3. *Given $\theta \in (0,\pi]$, there exists $\kappa_\theta \geq 0$ so that for all $\theta_0 \in [\theta,\pi]$ the map p_{θ_0} is a $(1,\kappa_\theta)$-quasi-isometric embedding.*

Proof. We show that it suffices to take

$$\kappa_\theta = 2\cdot \operatorname{arccosh}\left(\frac{1}{\sin\left(\frac{\theta}{2}\right)}\right).$$

Indeed, suppose that $a,b \in \mathbb{H}^2$, let $\overline{a} = p_\theta(a)$ and $\overline{b} = p_\theta(b)$ and consider the image of the geodesic segment $[a,b]$ in $p_\theta(\mathbb{H}^2)$. If the sign of the x-coordinates of a and b are the same, then $[a,b]$ maps to a geodesic segment in $\mathbb{H}^3$ and $d_{\mathbb{H}^3}(\overline{a},\overline{b}) = d_{\mathbb{H}^2}(a,b)$ in this case.

Suppose then that the signs of the x-coordinates of a and b are different, and let $c \in [a,b]$ have x-coordinate 0. Let $\overline{c} = p_\theta(c)$. Then $p_\theta([a,b])$ consists of two geodesic segments $[\overline{a},\overline{c}]$ and $[\overline{c},\overline{b}]$ meeting at some angle $\eta \geq \theta$.

Consider the geodesic triangle Δ in $\mathbb{H}^3$ with vertices $\overline{a},\overline{b},\overline{c}$, and let e be the distance from $\overline{c}$ to the geodesic $[\overline{a},\overline{b}]$. The shortest geodesic from $\overline{c}$ to $[\overline{a},\overline{b}]$ cuts Δ into two right-angled hyperbolic triangles, one of which has angle at $\overline{c}$ at least $\frac{\theta}{2}$. We thus have a hyperbolic triangle with side lengths e_1,e_2,e_3, say, where the angle opposite e_3 is $\frac{\pi}{2}$, and the angle opposite e_2 is at $\overline{c}$ and is $E_2 \geq \frac{\theta}{2}$. Let E_1 be the angle opposite the side of length e_1.

The second hyperbolic law of cosines says

$$\cos(E_1) = -\cos(E_2)\cos\left(\frac{\pi}{2}\right) + \sin(E_2)\sin\left(\frac{\pi}{2}\right)\cosh(e_1),$$

so

$$\cosh(e_1) = \frac{\cos(E_1)}{\sin(E_2)} \le \frac{1}{\sin\left(\frac{\theta}{2}\right)}.$$

Let $d_1 = d_{\mathbb{H}^3}(\overline{a},\overline{c})$ and $d_2 = d_{\mathbb{H}^3}(\overline{c},\overline{b})$. Observe that $d_{\mathbb{H}^2}(a,b) = d_1 + d_2$. It is clear that

$$d_1 + d_2 - 2\operatorname{arccosh}\left(\frac{1}{\sin\left(\frac{\theta}{2}\right)}\right) \le d_{\mathbb{H}^3}(\overline{a},\overline{b}) \le d_1 + d_2,$$

and the result follows. □

3. Kahn–Wright surfaces

From this section until the end of Section 5, let N be a (noncompact) finite-volume hyperbolic 3-manifold. We remark that our proof of Theorem 1.1 works in the compact hyperbolic setting also without changes and is simpler in some ways than Sun's in this case.

The set of closed geodesics in $N = \mathbb{H}^3/\Gamma$ is in 1-to-1 correspondence with the set of conjugacy classes of loxodromic elements in Γ. For a closed geodesic α in N (with corresponding conjugacy class $[\gamma] \subset \Gamma$), let $\ell(\alpha)$ denote the length of α (the translation length of γ) and $\theta(\alpha)$ the holonomy class of α (the rotation angle of γ around its axis).

3.1. Pre-good curves

Later in the section, we give a brief discussion of the construction of surfaces due to Kahn and Wright in [15]. However, we first give a lemma which proves the existence of certain well-behaved geodesics whose n^{th} powers will become part of the Kahn–Wright surface. In the next section, we take a cover of the Kahn–Wright surface, cut along a lift of this n^{th} power, and then quotient by the n^{th} root of the two resulting boundary curves, to form a complex X_n. This construction is similar to Sun's in [20] and forms the basis of the proof of Theorem 1.1 in the finite-volume hyperbolic case. The following is an analogue in the finite-volume case of Sun's [20, Lemma 2.9]. In order to use this geodesic in Kahn and Wright's construction, it is important to control its *height*, a measure of how far it goes into a cusp neighborhood. See [15, §3] for the definition of height in the following statement.

Lemma 3.1. *For $n \in \mathbb{N}, \epsilon > 0$, $\mu > \frac{1}{2}$, there exists R_0 so that for all $R > R_0$ there exists a geodesic α_0 in N of height at most $\mu \log R$ such that $\left|\ell(\alpha_0) - \frac{2R}{n}\right| < \frac{\epsilon}{n}$ and $\left|\theta(\alpha_0) - \frac{2\pi}{n}\right| < \frac{\epsilon}{n}$.*

Proof. For a closed subset Ω of $\mathrm{SO}(2)$ and $T > 0$, let

$$\mathcal{G}(T,\Omega) = \{\alpha : \alpha \text{ is a closed geodesic in } N, \ell(\alpha) \le T, \theta(\alpha) \in \Omega\}.$$

As noted in [15, §3.1], an application of the Margulis argument shows that

$$\#\mathcal{G}(T,\Omega) \sim \frac{e^{2T}}{2T}\|\Omega\| \text{ as } T \to \infty \tag{1}$$

which in this case follows, for example, from [18, Theorem 1.1] by setting $\varphi := 1_\Omega$ the indicator function on $\mathrm{SO}(2)$ (see also [9]).

Considering geodesics $\alpha \in \mathcal{G}(2R/n+\epsilon/n,\Omega)\setminus\mathcal{G}(2R/n-\epsilon/n,\Omega)$, where Ω is the interval $(\frac{2\pi}{n}-\frac{\epsilon}{n},\frac{2\pi}{n}+\frac{\epsilon}{n})$, we have

$$\#\left\{\alpha : \left|\ell(\alpha_0)-\frac{2R}{n}\right| < \frac{\epsilon}{n} \text{ and } \left|\theta(\alpha_0)-\frac{2\pi}{n}\right| < \frac{\epsilon}{n}\right\} \sim \mu_\epsilon \frac{e^{4R/n}}{4R}. \tag{2}$$

The arguments in the proof of [15, Lemma 3.1] apply to show that, as R grows, the proportion of those α with height larger than $\mu \log R$ shrinks since $\mu > \frac{1}{2}$. In particular, for sufficiently large R, one can find α_0 as needed. □

Note that α_0 may be chosen to be primitive. In the language of Kahn and Wright, α_0^n is an (R,ϵ)-*good curve*.

Definition 3.2. Fix $n \in \mathbb{N}$ and also R,ϵ. An (R,ϵ,n)-*pre-good curve* in N is a geodesic α_0 satisfying the conclusion of Lemma 3.1 for $\mu = 1$.

We remark that Kahn and Wright allow curves to have height at most $50\log(R)$ before needing to be 'cut-off', so certainly curves of height at most $\log R$ are fine. Lemma 3.1 asserts that, for fixed n and ϵ, for large enough R, there exists an (R,ϵ,n)-pre-good curve (in fact there are many).

3.2. The construction of Kahn and Wright

In [15], Kahn and Wright build certain quasi-Fuchsian immersed surfaces in N out of pieces called *good pants* and *good hamster wheels*. Each good pant and good hamster wheel is immersed in N and has geodesic boundary components, which are referred to as *cuffs*.

The construction in [15] depends on choices of parameters R (sufficiently large) and $\epsilon > 0$ (sufficiently small). A pant is a sphere with three holes, and a hamster wheel is a sphere with $R+2$ holes, 2 of which are special. Pants and hamster wheels are good if all cuffs have complex length within ϵ of $2R$ and perfect if they are exactly $2R$.

We postpone for now the choice of the parameters R and ϵ in order to discuss the construction. Kahn and Wright also specify another pair of parameters called 'cutoff heights', and the purpose of Lemma 3.1 above is to ensure that we can find an α_0 whose height stays below the cutoff heights and whose n^{th} power is a good curve.

Suppose that we find a curve α_0 as in Lemma 3.1 so that $\alpha = \alpha_0^n$ is an (R,ϵ)-good curve and so that the height of α_0 (and hence α) is at most $\log R$. Then Kahn and Wright build a surface S out of good pants and good hamster wheels, and α appears as a cuff on at least one (in fact many) of these pieces.

Let $H_R = \{(x,y) \in \mathbb{H}^2 \mid x \geq 0\}$ and $H_L = \{(x,y) \in \mathbb{H}^2 \mid x \leq 0\}$, and let $m = \{(x,y) \in \mathbb{H}^2 \mid x = 0\} = H_R \cap H_L$.

Consider $\alpha_0 \colon \mathbb{S}^1 \to N$ as a map from the circle to N parametrized proportional to arc length, and let $u = \alpha_0(1)$. Let $\phi_n \colon \mathbb{S}^1 \to \mathbb{S}^1$ be the connected n-fold covering map, and let $\alpha \colon \mathbb{S}^1 \to N$ be the composition $\alpha_0 \circ \phi_n$. Suppose $f \colon S_0 \looparrowright N$ is a Kahn–Wright surface

https://doi.org/10.1017/S1474748022000329 Published online by Cambridge University Press

and that there is a map $C\colon \mathbb{S}^1 \to S_0$ so $\alpha = f \circ C$. Let $\phi_n^{-1}(u) = \{u_1,\ldots,u_n\}$, and note that there are n different points $\{a_1,\ldots,a_n\}$ on S_0, so $f(a_i) = \alpha(u_i)$ for each i.

Choose a basepoint $b \in \mathbb{H}^3$, and let $\pi\colon (\mathbb{H}^3,b) \to (N,u)$ be the based universal covering map. Fix a basepoint $c \in \mathbb{H}^2$, and for each i, let $\tau_i\colon (\mathbb{H}^2,c) \to (S_0,a_i)$ be a based universal cover so that $\tau_i(m) = C(\mathbb{S}^1)$.

The map f elevates to n distinct (based) maps:

$$\widetilde{f}_i\colon (\mathbb{H}^2,c) \to (\mathbb{H}^3,b)$$

so that for each i we have $\pi \circ \widetilde{f}_i = f \circ \tau_i$.

Now, for a pair $i \neq j$ from $\{1,\ldots,n\}$, we have $\widetilde{f}_i(m) = \widetilde{f}_j(m)$. Thus, we can take the two maps $\widetilde{f}_i|_{H_R}$ and $\widetilde{f}_j|_{H_R}$ and glue them together via an orientation-preserving isometry along the boundary to get a continuous map $\widetilde{f}_{i,j}^{H_R}\colon \mathbb{H}^2 \to \mathbb{H}^3$, and similarly for the two maps restricted to H_L to get a continuous map $\widetilde{f}_{i,j}^{H_L}$.

Kahn and Wright prove that, for appropriate choices of parameters, their surface, built as an assembly of good pants and good hamster wheels, is close to an assembly of perfect pants and perfect hamster wheels and that the map which takes the 'good' assembly to the 'perfect' assembly is *compliant* (see [15, §A.5]), which in particular means that it takes cuffs to cuffs. For a perfect assembly with cuff α, the construction analogous to the $\widetilde{f}_{i,j}$ leads to pairs of totally geodesic half-planes glued along their boundary geodesic, namely to a map p_θ for some θ. Thus, the map that takes the good assembly to the perfect assembly induces a map between $\widetilde{f}_{i,j}^{H_R}\colon \mathbb{H}^2 \to \mathbb{H}^3$ and some map $p_\theta\colon \mathbb{H}^2 \to \mathbb{H}^3$, and this map takes m to the pleating locus for p_θ.

Our first task is to bound θ away from 0, and our second is to show that the two maps are close. The sense in which they are close will be that of [15, p. 554]—being of ϵ_0*-bounded distortion to distance* D for appropriate choice of ϵ_0 and D.

Denote the angles of the maps p_θ induced by i, j and H_R by $\theta(i,j,H_R)$ and for i, j and H_L by $\theta(i,j,H_L)$.

The following is a summary of the above discussion and also of [15, Theorem A.18]. In the following statement, R_0 is the constant from the statement of Lemma 3.1 (with values n, ϵ_0 and $\mu = 1$, respectively).

Theorem 3.3. *Fix $n \in \mathbb{N}$. For all D there exist C, ϵ_0 and $R_1 > R_0$ so that for all $\epsilon \in (0,\epsilon_0)$ and all $R > R_1$ and any (n,R,ϵ)-pre-good curve α_0 there exists a Kahn–Wright surface $f\colon S_0 \looparrowright N$ containing $\alpha = \alpha_0^n$ as a cuff, constructed as an (R,ϵ)-good assembly.*

For each $i,j \in \{1,\ldots,n\}$ with $i \neq j$, we have the following:

(1) *For any point v which lies within D of m in $\mathbb{H}^2$, we have*

$$d\left(\widetilde{f}_{i,j}^{H_L}(v), p_{\theta(i,j,H_L)}(v)\right), d\left(\widetilde{f}_{i,j}^{H_R}(v), p_{\theta(i,j,H_R)}(v)\right) < C\epsilon$$

and

(2) *For any point $z \in \mathbb{H}^2$ lying at distance at least D from m, the maps $\widetilde{f}_{i,j}^{H_L}$ and $\widetilde{f}_{i,j}^{H_R}$ restricted to the ball of radius D about z are $(1+C\epsilon, C\epsilon)$-quasi-isometric embeddings.*

Moreover, we have

$$\theta(i,j,H_R),\theta(i,j,H_L) \in \left(\frac{\pi}{n},\pi\right).$$

The following is an easy consequence of Theorem 3.3 and Lemma 2.3. In the following statement ϵ_D is the constant from Theorem 3.3 applied to n and D, R_0 is the constant from Lemma 3.1 applied with choices n, ϵ_D and $\mu = 1$, and R_1 is the constant then obtained from Theorem 3.3.

Corollary 3.4. *Fix $n \in \mathbb{N}$. There exist λ,κ so that for any D there exist $R_D > R_1, R_0$ and $\epsilon_D > 0$ so that for any α_0 and $f\colon S_0 \looparrowright N$ as in Theorem 3.3 with $R > R_D$ and any $\epsilon \in (0,\epsilon_D)$ the maps $\widetilde{f}_{i,j}^{H_R}$ and $\widetilde{f}_{i,j}^{H_L}$ are D-local (λ,κ)-quasi-isometric embedding.*

Now, choose D,λ_1,κ_1 so that any D–local (λ,κ)-quasi-isometric embedding from $\mathbb{H}^2$ to $\mathbb{H}^3$ is a global (λ_1,κ_1)-quasi-isometric embedding (see Proposition 2.2). This D then gives R_D and ϵ_D as above.

Lemma 3.1 proves that there is an (n,R_D,ϵ_D)-pre-good curve α_0, and the construction from [15] proves that there is an $f\colon S_0 \to N$ with $\alpha = \alpha_0^n$ as a cuff satisfying the conclusions of Theorem 3.3 and Corollary 3.4, with $R = R_D$.

We fix this map $f\colon S_0 \looparrowright N$, along with n, D, R_D, ϵ_D, α_0, $\alpha = \alpha_0^n$, κ and ϵ as chosen above for the next two sections.

4. The space X_n

By standard separability properties of surface groups, we may find a cover $S \to S_0$ to which α lifts as a nonseparating simple closed curve and so that:

(1) The injectivity radius of S is at least $\max\{2D,\lambda_1\kappa_1\}$, and
(2) The lift of α to S is contained in an embedded collar of width at least $\max\{2D,\lambda_1\kappa_1\}$.

Given the surface S, we build a space X_n which immerses into N, exactly as in [20]. Passing from S_0 to S before constructing X_n makes the proof that X_n is π_1-injective with quasi-convex image much simpler than Sun's proof from [20, §4]. Let C denote the image of α in S, and let $\phi_C^n\colon C \to \mathbb{S}^1$ be an n-to-1 covering map, and let $\tau_C\colon C \to C$ be a deck transformation. We may choose ϕ_C^n so that τ is an isometry.

Definition 4.1. The space $X_n(S,C)$ is defined by cutting S along C to get a surface S_1 with two boundary components, denoted C_1 and C_2, and taking the quotient of S_1 by the relation generated by $c \sim \tau_{C_i}(c)$ for $c \in C_i$ and $i = 1,2$.

Suppose that S is equipped with a hyperbolic metric, and consider the induced metric on S_1. Since the maps τ_{C_i} are isometries, there is a natural induced quotient metric on $X_n(S,C)$, which is locally isometric to $\mathbb{H}^2$ away from the images of the C_i.

The following result is clear from the construction of S from S_0.

Lemma 4.2. *The injectivity radius of X_n is at least* $\max\{2D,\lambda_1\kappa_1\}$.

https://doi.org/10.1017/S1474748022000329 Published online by Cambridge University Press

Let S_1 be the surface obtained from S by cutting along C, and let C_1, C_2 be the boundary components of S_1. Let $q\colon S_1 \to X_n$ be the defining quotient map, and let $\overline{C_i} = q(C_i)$ for $i = 1,2$.

Because in S the curve C has an embedded collar of width at least $2D$, for any $i,j \in \{1,2\}$, any two distinct elevations of $\overline{C_i}$ and $\overline{C_j}$ to $\widetilde{X_n}$ are at distance at least $4D$ from each other.

Definition 4.3. Suppose that $\mathcal{A} = \{Z_1, \ldots, Z_m\}$ is a finite collection of metric spaces and that $k > 0$. A metric space Z is k-modeled on $\mathcal{A}$ if for every $z \in Z$ there is an i so that the ball of radius k about z is isometric to a ball in Z_i.

Recall $H_R = \{(x,y) \mid x \geq 0, y > 0\}$ is the (closed) half-hyperbolic plane (in the upper half-space model). Let W_n be the space obtained from n copies of H_R glued along the boundary geodesics (by an isometry).

Lemma 4.4. *The space $\widetilde{X_n}$ is D-modeled on W_n.*

Proof. Let $x \in \widetilde{X_n}$, and consider the covering map $\pi\colon \widetilde{X_n} \to X_n$.

Case 1: $d(\pi(x), \{\overline{C_1}, \overline{C_2}\}) \leq D$.

In this case, in $\widetilde{X_n}$ there is a unique elevation of some $\overline{C_i}$ which lies within D of x. Let y be a point in this elevation so $d(x,y) \leq D$. Then $B_D(x) \subseteq B_{2D}(y)$, and $B_{2D}(y)$ is isometric to a ball of radius $2D$ in W_n.

Case 2: $d(\pi(x), \{\overline{C_1}, \overline{C_2}\}) > D$.

In this case, there is no elevation of either $\overline{C_i}$ which lies within D of x, and $B_D(x)$ is isometric to a ball of radius D in $\mathbb{H}^2$ (and so in W_n). □

By construction, the immersion $f_1\colon S \to N$ obtained from composing the covering map $S \to S_0$ with $f\colon S_0 \looparrowright N$ yields an immersion $g\colon X_n \looparrowright N$. Let $\widetilde{g}\colon \widetilde{X_n} \to \mathbb{H}^3$ be the induced map on universal covers.

Two points x,y in $\widetilde{X_n}$ at distance at most D either lie in an isometrically embedded copy of a half-space from $\mathbb{H}^2$,= or else in two different 'sheets' of a copy of W_n. In either case, it follows immediately from Corollary 3.4 that $\widetilde{g}$ is a D–local (λ,κ)-quasi-isometric embedding. Thus,

Theorem 4.5. *The map $\widetilde{g}\colon \widetilde{X_n} \to \mathbb{H}^3$ is a D-local (λ,κ)-quasi-isometric embedding and hence is a (global) (λ_1,κ_1)-quasi-isometric embedding.*

In particular, since the injectivity radius of X_n is at least $\lambda_1\kappa_1$ the map g is π_1-injective, and $g_*(\pi_1(X_n))$ is relatively quasi-convex in $\pi_1(N)$. Moreover, $g_*(\pi_1(X_n))$ does not intersect any (conjugate of) the cusp subgroups of $\pi_1(N)$.

5. Virtual retractions and the proof of Theorem 1.1 in the hyperbolic case

In this section, we prove Theorem 1.1 in case of a (noncompact) finite-volume hyperbolic 3-manifold N. Let $\Gamma = \pi_1(N)$. By [22, Theorem 17.14], $\pi_1(N)$ is the fundamental group of a compact virtually special cube complex X. Let $\Gamma' \leq \Gamma$ be a finite-index subgroup so that the cover of X corresponding to Γ' is special, and let N_1 be the cover of N corresponding to Γ'. As in Section 4, construct an immersion $g\colon X_n \to N_1$. Note that $(\Gamma', \mathcal{P})$ is relatively hyperbolic, where $\mathcal{P}$ consists of the (abelian) cusp subgroups.

Let $H = g_*(\pi_1(X_n)) \leq \pi_1(N_1) = \Gamma'$. The subgroup H is relatively quasi-convex in Γ', and so by [12, Corollary 6.7] (we use the formulation as in [19, Theorem 6.3]) that H is a virtual retract of Γ'. Let Γ'' be a finite-index subgroup of Γ' which retracts onto H. Let N_2 be the finite cover of N_1 corresponding to Γ''. As in [20, Proposition 3.7], we have the induced maps on homology:

$$H_1(X_n;\mathbb{Z}) \xrightarrow{g_*} H_1(N_2;\mathbb{Z}) \xrightarrow{r_*} H_1(X_n;\mathbb{Z}).$$

Therefore, since $r \circ g_* = id_H$, $H_1(X_n;\mathbb{Z}) = \mathbb{Z}^{2\cdot \text{genus}(S)+1} \oplus \mathbb{Z}/n\mathbb{Z}$ is a direct factor of $H_1(N_2;\mathbb{Z})$. In particular, $\mathbb{Z}/n\mathbb{Z}$ is a direct factor of $H_1(N_2;\mathbb{Z})$, and this proves the hyperbolic case of Theorem 1.1 in the case that A is finite cyclic.

Given a finite abelian group A, induction on the rank k of A also works as in [20, Proposition 3.9] as follows. Let $A = A_1 \oplus \mathbb{Z}/n_{k+1}\mathbb{Z}$, where $A_1 = \oplus_{i=1}^k \mathbb{Z}/n_i\mathbb{Z}$. Suppose by induction that $H_1 \leq \Gamma'$ is a relatively quasi-convex free product of images of $\pi_1(X_{n_i})$ (for $i = 1, \dots, k$) and that $H_2 = (g_{k+1})_*\left(\pi_1(X_{n_{k+1}})\right) \leq \Gamma'$. Choose any $\gamma \in \Gamma'$ whose fixed points in $\partial\mathbb{H}^3$ are disjoint from both limit sets $\Lambda(H_1)$ and $\Lambda(H_2)$. Then after conjugating H_2 by some sufficiently high power γ^m, the first Klein–Maskit combination theorem [17] applies (note that by [11, Corollary 1.3] a subgroup is relatively quasi-convex if and only if it is geometrically finite). Since H_2 is relatively quasi-convex, the free product $H_1 * \gamma^m H_2 \gamma^{-m}$ is also a relatively quasi-convex subgroup of Γ' isomorphic to the abstract group $H_1 * H_2$. The proof of the hyperbolic case of Theorem 1.1 for a general finite abelian A then follows exactly as in the case of a finite cyclic A above.

6. Nonhyperbolic manifolds

We now prove Theorem 1.1 in general. To that end, suppose that M is an irreducible 3-manifold which has at least one hyperbolic piece in its geometric decomposition and that A is a finite abelian group. By [19, Theorem 1.1], there exists a CAT(0) cube complex X equipped with a free $\pi_1(M)$–action so that there are finitely many orbits of hyperplanes and so that $\pi_1(M)\backslash X$ has a finite special cover. Let $\Gamma_1 \leq \pi(M)$ be a finite-index subgroup corresponding to the finite special cover of $\pi_1(M)\backslash X$, and let M_1 be the finite cover of M corresponding to Γ_1. Let M_h be a hyperbolic piece in the geometric decomposition of M_1, and let $\Gamma_h := \pi_1(M_h) \leq \pi_1(M_1)$ (basepoints/conjugacy classes are not important here). According to the construction in the previous sections, there is a relatively quasi-convex subgroup H of Γ_h so that A is a direct factor of $H_1(H;\mathbb{Z})$. Theorem 1.1 then follows immediately from the following result. This result is presumably known to the experts, but we were unable to find it in the literature.

Proposition 6.1. *There is a finite-index subgroup $\Gamma_0 \leq \pi_1(M_1)$ so that $H \leq \Gamma_0$ and H is a retract of Γ_0.*

Proof. Let $\widetilde{M_h} \leq \widetilde{M}$ be the (Γ_h–invariant) universal cover of M_h inside the universal cover of M. The space X is built via a *wallspace* construction on $\widetilde{M}$ as in [13]. As in [13,

§3.4], we can associate to $\widetilde{M_h}$ a *hemiwallspace* consisting of those half-spaces in $\widetilde{M}$ which intersect $\widetilde{M_h}$ (see [13, Example 3.20]. This builds a Γ_h-invariant convex subcomplex X_h of X by [13, Lemma 3.24].

By [19, Theorem 2.1], the surfaces of the cubulation intersecting M_h all intersect M_h in a geometrically finite surface. Therefore, by [13, Theorem 7.10], the Γ_h–action on X_h is (free and) *co-sparse* (see [21, Definition 7.1]).

According to [21, Theorem 7.2] inside of X_h there is an H-invariant convex subcomplex Z upon which H acts co-sparsely. In fact, since H does not intersect any of the parabolic subgroups of Γ_h, the subcomplex Z found in [21, Theorem 7.2] is H-cocompact (this follow immediately from the proof).

Since $\Gamma_1 \backslash X$ is special, it follows from [12, Corollary 6.7] (we use the formulation as in [19, Theorem 6.3]) that H is a virtual retract of Γ_1. □

Acknowledgements. The first author is supported in part by NSF grant DMS 1803094 and by MSRI. The second author is supported in part by NSF grant DMS 1904913.

Competing Interests. None.

References

[1] I. Agol, '*Systoles of hyperbolic* 4-manifolds', Preprint, 2006, arXiv:0612290.

[2] I. Agol, 'The virtual Haken conjecture', *Doc. Math.* **18** (2013) 1045–1087, With an appendix by Agol, Daniel Groves, and Jason Manning.

[3] N. Bergeron, F. Haglund and D. T. Wise, 'Hyperplane sections in arithmetic hyperbolic manifolds', *J. Lond. Math. Soc. (2)* **83**(2) (2011), 431–448.

[4] M. V. Belolipetsky and S. A. Thomson, 'Systoles of hyperbolic manifolds', *Algebr. Geom. Topol.* **11**(3) (2011), 1455–1469.

[5] M. Coornaert, T. Delzant and A. Papadopoulos, *Géométrie et théorie des groupes: Les groupes hyperboliques de Gromov*, Lecture Notes in Mathematics, Vol. **1441** (Springer-Verlag, Berlin, 1990).

[6] D. Cooper and D. Futer, 'Ubiquitous quasi-Fuchsian surfaces in cusped hyperbolic 3-manifolds', *Geom. Topol.* **23**(1) (2019), 241–298.

[7] S. Friedl and G. Herrmann, 'Torsion in the homology of finite covers of 3-manifolds', Preprint, 2017, arXiv:1710.08983.

[8] M. Gromov and I. Piatetski-Shapiro, 'Nonarithmetic groups in Lobachevsky spaces', *Inst. Hautes Études Sci. Publ. Math.* (**66**) (1988), 93–103.

[9] R. Gangolli and G. Warner, 'Zeta functions of Selberg's type for some noncompact quotients of symmetric spaces of rank one', *Nagoya Math. J.*, **78** (1980), 1–44.

[10] A. Hadari, 'Every infinite order mapping class has an infinite order action on the homology of some finite cover', Preprint, 2015, arXiv:1508.01555.

[11] G. Christopher Hruska, 'Relative hyperbolicity and relative quasiconvexity for countable groups', *Algebr. Geom. Topol.* **10**(3) (2010), 1807–1856.

[12] F. Haglund and D. T. Wise, 'Special cube complexes', *Geom. Funct. Anal.* **17**(5) (2008), 1551–1620.

[13] G. C. Hruska and D. T. Wise, 'Finiteness properties of cubulated groups', *Compos. Math.* **150**(3) (2014), 453–506.

[14] J. Kahn and V. Markovic, 'Immersing almost geodesic surfaces in a closed hyperbolic three manifold', *Ann. of Math. (2)* **175**(3) (2012), 1127–1190.
[15] J. Kahn and A. Wright, 'Nearly Fuchsian surface subgroups of finite covolume Kleinian groups', *Duke Math. J.* **170**(3) (2021), 503–573.
[16] Y. Liu, 'Virtual homological spectral radii for automorphisms of surfaces', *J. Amer. Math. Soc.* **33**(4) (2020), 1167–1227.
[17] B. Maskit, 'On Klein's combination theorem'. *IV', Trans. Amer. Math. Soc.* **336**(1) (1993), 265–294.
[18] G. Margulis, A. Mohammadi and H. Oh, 'Closed geodesics and holonomies for Kleinian manifolds', *Geom. Funct. Anal.* **24**(5) (2014), 1608–1636.
[19] P. Przytycki and D. T. Wise, 'Mixed 3-manifolds are virtually special', *J. Amer. Math. Soc.* **31**(2) (2018), 319–347.
[20] H. Sun, 'Virtual homological torsion of closed hyperbolic 3-manifolds', *J. Differential Geom.* **100**(3) (2015), 547–583.
[21] M. Sageev and D. T. Wise, 'Cores for quasiconvex actions', *Proc. Amer. Math. Soc.* **143**(7) (2015), 2731–2741.
[22] D. T. Wise, *The Structure of Groups with a Quasiconvex Hierarchy*, *Annals of Mathematics Studies*, Vol. **209** (Princeton University Press, Princeton, NJ, 2021).