

QUASI-ISOMETRIC CLASSIFICATION OF NON-GEOMETRIC 3-MANIFOLD GROUPS

JASON A. BEHRSTOCK AND WALTER D. NEUMANN

ABSTRACT. We describe the quasi-isometric classification of fundamental groups of irreducible non-geometric 3-manifolds which do not have “too many” arithmetic hyperbolic geometric components, thus completing the quasi-isometric classification of 3-manifold groups in all but a few exceptional cases.

1. INTRODUCTION

In [1] we discussed the quasi-isometry classification of fundamental groups of 3-manifolds with possibly empty toroidal boundary (which coincides with the bilipschitz classification of the universal covers of these three-manifolds). This classification reduces easily to the case of irreducible manifolds. Moreover, no generality is lost by considering only orientable manifolds. So from now on we use the term 3-manifolds to refer to those we need consider, namely the compact connected orientable irreducible 3-manifolds of zero Euler characteristic (i.e., with boundary consisting only of tori) since these are the orientable manifolds which, by Perelman’s Geometrization Theorem [11, 12, 13], decompose along tori and Klein bottles into geometric pieces (this decomposition removes the boundary tori and a family of embedded tori and Klein bottles, so the pieces of the decomposition are without boundary). The minimal such decomposition is what is called the *geometric decomposition*.

We described (loc. cit.) the classification for geometric 3-manifolds as well as for non-geometric 3-manifolds with no hyperbolic pieces in their geometric decomposition (i.e., graph-manifolds). For geometric manifolds this was a summary of work of others; our contribution was in the non-geometric case. In this paper we extend to allow hyperbolic pieces. However, our results are still not quite complete: at present we exclude manifolds with “too many” arithmetic hyperbolic pieces¹.

In the bulk of this paper we restrict to non-geometric manifolds, all of whose geometric components are hyperbolic and at least one of which is non-arithmetic (we will call these \mathcal{N} -manifolds for short). In the final section 8 we extend to the case where Seifert fibered pieces are also allowed.

The classification for graph-manifolds in [1] was in terms of finite labeled graphs; the labeling consisted of a color black or white on each vertex and the classification was in terms of an equivalence relation called *bisimilarity*. Each bisimilarity class

2000 *Mathematics Subject Classification*. Primary 20F65; Secondary 57N10, 20F36.

Key words and phrases. three-manifold, quasi-isometry, commensurability.

Research supported under PSC-CUNY grant no. 63089-00 41.

Research supported under NSF grant no. DMS-0604524.

¹The essential remaining case to address is when *all* pieces are arithmetic hyperbolic. These behave rather differently from the other cases—more like arithmetic hyperbolic manifolds.

contained a unique minimal two-colored graph, so these minimal two-colored graphs served as classifying objects. For \mathcal{N} -manifolds the classification is again in terms of finite labeled graphs and a “bisimilarity” relation, with a unique minimal labeled graph in each bisimilarity class which can serve as the classifying object. The labeling is more complex: each vertex is labeled by the isomorphism type of a hyperbolic orbifold² and each edge is labeled by a linear isomorphism between certain 2-dimensional \mathbb{Q} -vectorspaces. We will call these graphs \mathcal{N} -graphs if at least one orbifold label is non-arithmetic (\mathcal{H} -graphs without this condition). The graphs that classify \mathcal{N} -manifolds are the \mathcal{N} -graphs that are *minimal* and *balanced*. All these concepts will be defined in Section 3.

Finally, when both Seifert fibered and hyperbolic pieces occur the classifying graphs are a hybrid of the two-colored graphs and the \mathcal{H} -graphs, so we call them \mathcal{HS} -graphs. In the case that every component of the manifold obtained by removing all Seifert fibered pieces is \mathcal{N} we call the hybrid graphs \mathcal{NS} -graphs and we call irreducible non-geometric manifolds of this type *good*. For example, 3-manifolds which contain a hyperbolic piece but no arithmetic hyperbolic piece provide a large family of good manifolds.

The following three theorems summarize our main results. The second two theorems complement the first, by making the classification effective, and then relating the quasi-isometric and commensurability classification for some \mathcal{N} -manifolds.

Classification Theorem. *Each good 3-manifold has an associated \mathcal{NS} -graph and two such manifolds have quasi-isometric fundamental groups (in fact, bilipschitz equivalent universal covers) if and only if their \mathcal{NS} -graphs are bisimilar, or equivalently, the associated minimal \mathcal{NS} -graphs are isomorphic.*

Realization Theorem. *The minimal \mathcal{NS} -graph associated to a good 3-manifold is balanced. Any minimal \mathcal{NS} -graph which is balanced is the minimal \mathcal{NS} -graph for some quasi-isometry class of good 3-manifold groups.*

Commensurability Theorem. *If two \mathcal{N} -manifolds have quasi-isometric fundamental groups and their common minimal \mathcal{N} -graph is a tree with manifold labels then they (and in particular, their fundamental groups) are commensurable.*

In the absence of the condition on existence of a non-arithmetic hyperbolic piece our results are much weaker. There is no longer a unique minimal graph in each bisimilarity class of \mathcal{HS} -graphs; if the graph is not \mathcal{NS} there are infinitely many. We still have that manifolds with bisimilar \mathcal{HS} -graphs have bilipschitz equivalent universal covers, but we do not know the converse (Remark 4.3). The Realization Theorem remains true: any bisimilarity class which contains a balanced \mathcal{HS} -graph, contains an \mathcal{HS} -graph of some 3-manifold (Section 6 and end of Section 8).

We do not know to what extent our restrictive conditions on the minimal \mathcal{N} -graph in the Commensurability Theorem are needed; whether “ \mathcal{N} -manifolds are quasi-isometric if and only if they are commensurable” holds in complete generality remains a very interesting open question. In the presence of Seifert fibered pieces commensurability is strictly more restrictive than quasi-isometry, c.f., [1, 9].

The converse part of the Realization theorem and the Commensurability Theorem both depend on the “Cusp Covering Conjecture” in dimension 3, which we explain

²“Hyperbolic orbifold” always means an orientable complete hyperbolic 3-orbifold of finite volume.

in Section 5. Since the original version of this paper, Dani Wise has pointed out that this conjecture follows from work in his preprint [15], and has now included a proof as Corollary 18.11 of [15], see also [16]. Although not relevant to the present results, we note that for dimension ≥ 4 the Cusp Covering Conjecture remains open.

A slightly surprising byproduct of this investigation are the *minimal orbifolds* of section 2, which play a role similar to commensurator quotients but exist also for cusped arithmetic hyperbolic orbifolds. Although their existence is easy to prove, they were new to us. Minimal orbifolds are precisely the orbifolds that can appear as vertex labels of minimal \mathcal{H} -graphs.

2. MINIMAL ORBIFOLDS

We consider only orientable manifolds and orbifolds. Let N be a hyperbolic 3-orbifold (not necessarily non-arithmetic) with at least one cusp. Then each cusp of N has a smallest cover by a toral cusp (one with toral cross section) and this cover has cyclic covering transformation group F_C of order 1, 2, 3, 4, or 6. We call this order the “orbifold degree of the cusp.”

Proposition 2.1. *Among the orbifolds N' covered by N having the same number of cusps as N with each cusp of N' having the same orbifold degree as the cusp of N that covers it, there is a unique one, N_0 , that is covered by all the others. We call this N_0 a minimal orbifold.*

More generally the same conclusion holds if for each cusp of N we specify a “target” in $\{1, 2, 3, 4, 6\}$ that is a multiple of the cusp orbifold degree and replace “having the same orbifold degree as the cusp of $N \dots$ ” by “having orbifold degree dividing the target degree for the cusp of $N \dots$ ” in the above.

Proof. We first “neuter” N by removing disjoint open horoball neighborhoods of the cusps to obtain a compact orbifold with boundary which we call N^0 . If we prove the proposition for N^0 , with boundary components interpreted as cusps, then it holds for N .

Let \tilde{N}^0 be the universal cover of N^0 . Any boundary component C of N^0 is isomorphic to the quotient of a euclidean plane \tilde{C} by the orbifold fundamental group $\pi_1^{orb}(C)$, which is an extension of a lattice \mathbb{Z}^2 by the cyclic group F_C of order 1, 2, 3, 4, or 6. We choose an oriented foliation of this plane by parallel straight lines and consider all images by covering transformations of this foliation on boundary planes of \tilde{N}^0 . Each boundary plane of \tilde{N}^0 that covers C will then have $|F_C|$ oriented foliations, related by the action of the cyclic group F_C . If the target degree n_C for the given cusp is a proper multiple of $|F_C|$, we also add the foliations obtained by rotating by multiples of $2\pi/n_C$.

We give the boundary planes of \tilde{N}^0 different labels according to which boundary component of N^0 they cover, and we construct foliations on them as above. So each boundary plane of \tilde{N}^0 carries a finite number (1, 2, 3, 4, or 6) of foliations and the collection of all these oriented foliations and the labels on boundary planes are invariant under the covering transformations of the covering $\tilde{N}^0 \rightarrow N^0$. Let G_N be the group of all orientation preserving isometries of \tilde{N}^0 that preserve the labels of the planes and preserve the collection of oriented foliations. Note that G_N does not depend on choices: The only relevant choices are the size of the horoballs removed when neutering and the direction of the foliation we first chose on a boundary plane of \tilde{N}^0 . If we change the size of the neutering and rotate the direction of the

foliation then the size of the neutering and direction of the image foliations at all boundary planes with the same label change the same amount, so the relevant data are still preserved by G_N . It is clear that G_N is discrete (this is true for any group of isometries of \mathbb{H}^3 which maps a set of at least three disjoint horoballs to itself). $N_0^0 := \tilde{N}^0/G_N$ is thus an oriented orbifold; it is clearly covered by N^0 , has the same number of boundary components as N^0 , and the boundary component covered by a boundary component C of N^0 has orbifold degree dividing the chosen target degree n_C . Moreover if N_1^0 is any other orbifold covered by N^0 and with the same properties, then $\tilde{N}_1^0 = \tilde{N}^0$ and the labelings and foliations on \tilde{N}_1^0 and \tilde{N}^0 can be chosen the same, so N_0^0 is also the minimal orbifold for N_1^0 . The proposition thus follows. \square

Note that a minimal orbifold may have a non-trivial isometry group (in contrast with commensurator quotients of non-arithmetic hyperbolic manifolds).

3. \mathcal{H} -GRAPHS AND \mathcal{N} -GRAPHS

The graphs we need will be finite, connected, undirected graphs; note that we allow multiple edges as well as self-loops. We take the viewpoint that an edge of an undirected graph consists of a pair of oppositely directed edges. The reversal of a directed edge e will be denoted \bar{e} and the initial and terminal vertices of an edge will be denoted ιe and $\tau e (= \iota \bar{e})$.

We will label the vertices of our graph by hyperbolic orbifolds so we first introduce some terminology for these. A horosphere section C of a cusp of a hyperbolic orbifold N will be called the *cuspid orbifold*. Although the position of C as a horosphere section of the cusp involves choice, as a flat 2-dimensional orbifold, C is canonically determined up to similarity by the cusp. We thus have one cuspid orbifold for each cusp of N .

Since a cuspid orbifold C is flat, away from the orbifold points its tangent space TC is independent of which point on C we choose, up to the action of the finite cyclic group F_C . TC naturally contains the maximal lattice $\mathbb{Z}^2 \subset \pi_1^{orb}(C)$, so it makes sense to talk of a linear isomorphism between two of these tangent spaces as being *rational*, i.e., given by a rational matrix with respect to oriented bases of the underlying integral lattices. Moreover, such a rational linear isomorphism will have a well-defined determinant.

Definition 3.1 (\mathcal{H} -graph). An \mathcal{H} -graph is a finite connected graph with the following data labeling its vertices and edges:

- (1) Each vertex v is labeled by a hyperbolic orbifold N_v plus a map $e \mapsto C_e$ from the set of directed edges e exiting that vertex to the set of cuspid orbifolds C_e of N_v . This map is injective, except that an edge which begins and ends at the same vertex may have $C_e = C_{\bar{e}}$.
- (2) The cuspid orbifolds C_e and $C_{\bar{e}}$ have the same orbifold degree.
- (3) Each directed edge e is labeled by a rational linear isomorphism $\ell_e: TC_e \rightarrow TC_{\bar{e}}$, with $\ell_{\bar{e}} = \ell_e^{-1}$. Moreover ℓ_e reverses orientation (so $\det(\ell_e) < 0$).
- (4) We only need ℓ_e up to right multiplication by elements of $F_e := F_{C_e}$; in other words, the relevant datum is really the coset $\ell_e F_e$ rather than ℓ_e itself. This necessitates:

- (5) ℓ_e conjugates the cyclic group F_e to the cyclic group $F_{\bar{e}}$, i.e., $\ell_e F_e = F_{\bar{e}} \ell_e$ (this holds automatically if the orbifold degree is ≤ 2 ; otherwise it is equivalent to saying that ℓ_e is a similarity for the euclidean structures).

Definition 3.2. We call an \mathcal{H} -graph *balanced* if the product of the determinants of the linear maps on edges around any closed directed path is equal to ± 1 .

We call the \mathcal{H} -graph *integral* if each ℓ_e is an integral linear isomorphism (i.e., an isomorphism of the underlying \mathbb{Z} -lattices). Integral clearly implies balanced.

Definition 3.3. An integral \mathcal{H} -graph contains precisely the information to specify how to glue neutered versions N_v^0 of the orbifolds N_v together along their boundary components to obtain a \mathcal{H} -orbifold M , called the *associated orbifold*. Conversely, any \mathcal{H} -orbifold M has an *associated integral \mathcal{H} -graph* $\Gamma(M)$, which encodes its decomposition into geometric pieces.

We next want to define morphisms of \mathcal{H} -graphs.

Definition 3.4 (Morphism). A *morphism* of \mathcal{H} -graphs, $\Gamma \rightarrow \Gamma'$, consists of the following data:

- (1) an abstract graph homomorphism $\phi: \Gamma \rightarrow \Gamma'$, and
- (2) for each vertex v of Γ , a covering map $\pi_v: N_v \rightarrow N_{\phi(v)}$ of the orbifolds labeling v and $\phi(v)$ which respects cusps (so for each departing edge e at v one has $\pi_v(C_e) = C_{\phi(e)}$), subject to the condition:
- (3) for each directed edge e of Γ the following diagram commutes

$$\begin{array}{ccc} TC_e & \xrightarrow{\ell_e} & TC_{\bar{e}} \\ \downarrow & & \downarrow \\ TC_{\phi(e)} & \xrightarrow{\ell_{\phi(e)}} & TC_{\phi(\bar{e})} \end{array}$$

Here the vertical arrows are the induced maps of tangent spaces and commutativity of the diagram is up to the indeterminacy of item (4) of Definition 3.1.

Morphisms compose in the obvious way, so a morphism is an *isomorphism* if and only if ϕ is a graph isomorphism and each π_v has degree 1.

Definition 3.5. Existence of a morphism between \mathcal{H} -graphs is a binary relation on the set of \mathcal{H} -graphs. It generates an equivalence relation which we call *bisimilarity*. An \mathcal{H} -graph is *minimal* if every morphism to an \mathcal{H} -graph is an isomorphism

Note that a morphism from an \mathcal{H} -graph to itself is an isomorphism, and thus if an \mathcal{H} -graph admits morphisms from every \mathcal{H} -graph in its bisimilarity class it is minimal.

Definition 3.6. If at least one of the vertex labels of an \mathcal{H} -graph is non-arithmetic, we call the graph an \mathcal{N} -graph.

Theorem 3.7. *Every bisimilarity class of \mathcal{N} -graphs contains a unique (up to isomorphism) minimal member, and every \mathcal{N} -graph in the bisimilarity class has a morphism to this unique minimal member.*

\mathcal{N} -graphs play an analogous role to the two-colored graphs that we used in [1] to classify graph-manifold groups up to quasi-isometry. For those graphs the morphisms were called “weak coverings”—they were color preserving open graph

homomorphisms (i.e., open as maps of 1-complexes). The term “bisimilarity” was coined in that context, since the concept is known to computer science under this name. We postpone the proof of Theorem 3.7 to the next section, since it follows naturally from the discussion there.

Theorem 3.8. *If $\Gamma \rightarrow \Gamma'$ is a morphism of \mathcal{H} -graphs and Γ is balanced, then so is Γ' . In particular, a bisimilarity class contains a balanced \mathcal{H} -graph if and only if it contains a minimal \mathcal{H} -graph which is balanced. Hence, for a balanced \mathcal{N} -graph the unique minimal graph is balanced.*

Proof. We will denote the negative determinant of the linear map labeling an edge e of an \mathcal{N} -graph by δ_e , so $\delta_e = \delta_{\bar{e}}^{-1} > 0$.

Assume Γ is balanced. We will show Γ' is balanced. We only need to show that the product of the δ_e 's along any simple directed cycle in Γ' is 1. Let $\mathcal{C} = (e'_0, e'_1, \dots, e'_{n-1})$, be such a cycle, so $\tau e'_i = \iota e'_{i+1}$ for each i (indices modulo n). Denote

$$(1) \quad D := \prod_{i=0}^{n-1} \delta_{e'_i},$$

so we need to prove that $D = 1$.

From now on we will restrict attention to one connected component Γ_0 of the full inverse image of \mathcal{C} under the map $\Gamma \rightarrow \Gamma'$.

Let e be an edge of Γ_0 which maps to an edge e' of \mathcal{C} (e' may be an e'_i or an \bar{e}'_i). The initial and terminal segments of this edge can each be associated with the cusp orbifold which covers the cusp orbifold corresponding to the 3-orbifold boundary associated to the initial or terminal segment of e' with covering degrees that we shall call d_e and $d_{\bar{e}}$ respectively. We claim that

$$(2) \quad d_e \delta_{e'} = \delta_e d_{\bar{e}}.$$

Indeed, if the cusp orbifolds are all tori, this equation (multiplied by -1) just represents products of determinants for the two ways of going around the commutative diagram in part (3) of Definition 3.4. In general, d_e and $d_{\bar{e}}$ are the determinants multiplied by the order of the cyclic group F_e , but this is the same factor for each of d_e and $d_{\bar{e}}$ so the equation remains correct.

Let $\tilde{\mathcal{C}} \rightarrow \mathcal{C}$ be the infinite cyclic cover of \mathcal{C} , and $\tilde{\Gamma}_0 \rightarrow \Gamma_0$ the pulled back infinite cyclic cover of Γ_0 . Define d_e or δ_e for an edge of either of these covers as the value for the image edge.

If (e_1, e_2, \dots, e_m) is a directed path in $\tilde{\Gamma}_0$ then the product of the equations (2) over all edges of this path gives $\prod_{i=1}^m d_{e_i} \delta_{\pi e_i} = \prod_{i=1}^m \delta_{e_i} d_{\bar{e}_i}$, so

$$\prod_{i=1}^m \frac{\delta_{\pi e_i}}{\delta_{e_i}} = \prod_{i=1}^m \frac{d_{\bar{e}_i}}{d_{e_i}}.$$

Since $\tilde{\Gamma}_0$ and $\tilde{\mathcal{C}}$ are both balanced, it follows that $\prod_{i=1}^m \frac{d_{\bar{e}_i}}{d_{e_i}}$ only depends on the start and end vertex of the path. Thus, if we fix a base vertex v_0 of $\tilde{\Gamma}_0$ and define

$$(3) \quad c(v) := \prod_{i=1}^m \frac{\delta_{\pi e_i}}{\delta_{e_i}} = \prod_{i=1}^m \frac{d_{\bar{e}_i}}{d_{e_i}} \quad \text{for any path from } v_0 \text{ to } v,$$

we get a well defined invariant of the vertices of $\tilde{\Gamma}_0$. For an edge e this invariant satisfies

$$\frac{c(\iota\bar{e})}{c(\iota e)} = \frac{c(\tau e)}{c(\tau\bar{e})} = \frac{d_{\bar{e}}}{d_e},$$

whence

$$\zeta(e) := \frac{d_e}{c(\iota e)}$$

is an invariant of the undirected edge: $\zeta(e) = \zeta(\bar{e})$.

Denote the map $\tilde{\Gamma}_0 \rightarrow \tilde{\mathcal{C}}$ by π . Number the vertices sequentially along $\tilde{\mathcal{C}}$ by integers $i \in \mathbb{Z}$, and for each vertex v of $\tilde{\Gamma}_0$ let $i(v) \in \mathbb{Z}$ be the index of $\pi(v)$. Let $h: \tilde{\Gamma}_0 \rightarrow \tilde{\Gamma}_0$ be the covering transformation: $h(v)$ is the vertex of $\tilde{\Gamma}_0$ with the same image in Γ_0 as v but with $i(h(v)) = i(v) + n$.

Note that the equations (1) and (3) and the fact that Γ_0 is balanced implies that $c(h(v)) = Dc(v)$ for any vertex v of $\tilde{\Gamma}_0$, so

$$(4) \quad \zeta(h(e)) = D^{-1}\zeta(e).$$

From now on consider the edges of $\tilde{\Gamma}_0$ directed only in the direction of increasing $i(v)$. Denote by $Z(j)$ the sum of $\zeta(e)$ over all edges with $i(\tau e) = j$.

The sum of the d_e 's over outgoing edges at a vertex v of $\tilde{\Gamma}_0$ equals the sum of $d_{\bar{e}}$'s over incoming edges at v , since each sum equals the degree of the covering map from the orbifold labeling v to the one labeling vertex $\pi(v)$ of $\tilde{\mathcal{C}}$. Since $\zeta(e) = d_e/c(v)$ for an outgoing edge and $\zeta(\bar{e}) = d_{\bar{e}}/c(v)$ for an incoming one, the sum of $\zeta(e)$ over outgoing edges at v equals the sum of $\zeta(e)$ for incoming ones. Thus $Z(j)$ is the sum of $\zeta(e)$ over all edges with $i(\iota e) = j$. These are the edges with $i(\tau e) = j + 1$, so $Z(j) = Z(j + 1)$. Thus $Z(j)$ is independent of j . Clearly $Z(j) > 0$. Equation (4) implies $Z(j + n) = D^{-1}Z(j)$, so $D = 1$, as was to be proved. (We are grateful to Don Zagier for help with this proof.) \square

We close this section with an observation promised in the introduction.

Proposition 3.9. *An orbifold can be a vertex label in a minimal \mathcal{H} -graph if and only if it is a minimal orbifold.*

Proof. We first prove ‘‘only if.’’ In condition (3) of the definition of a morphism, $\ell_{\phi(e)}$ is determined by ℓ_e since the vertical arrows are isomorphisms (but ℓ_e may not be determined by $\ell_{\phi(e)}$, since the indeterminacy $F_{\phi(e)}$ of $\ell_{\phi(e)}$ may be greater than the indeterminacy F_e of ℓ_e). Thus for any \mathcal{H} -graph there is an outgoing morphism to a new \mathcal{H} -graph for which the underlying graph homomorphism is an isomorphism and every orbifold vertex label is simply replaced in the new graph by the corresponding minimal orbifold; the new edge labels are then as just described.

For the converse, if N is a minimal orbifold, any \mathcal{H} -graph which is star-shaped, with N labeling the middle vertex and with one-cusp non-arithmetic commensurator quotients labeling the outer ones, is minimal. One needs a one-cusp commensurator quotient for each cusp degree in $\{1, 2, 3, 4, 6\}$ for this construction; these are not hard to find. \square

4. MINIMAL \mathcal{N} -GRAPHS CLASSIFY FOR QUASI-ISOMETRY

Let $M = M^3$ be a \mathcal{N} -manifold. For simplicity of exposition we first discuss the case that M has no arithmetic pieces, and then discuss the modifications needed when arithmetic pieces also occur.

M is then pasted together from pieces M_i , each of which is a neutered non-arithmetic hyperbolic manifold. We can choose neuterings in a consistent way, by choosing once and for all a neutering for the commensurator quotient in each commensurability class of non-arithmetic manifolds, and choosing each M_i to cover one of these “standard neutered commensurator quotients.”

The pasting identifies pairs of flat boundary tori with each other by affine maps which may not be isometries. To obtain a smooth metric on the result we glue a toral annulus $T^2 \times I$ between the two boundaries with a metric that interpolates between the flat metrics at the two ends in a standard way (if g_0 and g_1 are the flat product metrics induced on $T^2 \times I$ by the flat metric on its left and right ends we use $(1 - \rho(t))g_0 + \rho(t)g_1$ where $\rho: [0, 1] \rightarrow [0, 1]$ is some fixed smooth bijection with derivatives 0 at each end). This specifies the metric on M up to rigid translations of the gluing maps, so we get a compact family of different metrics on M .

The universal cover \widetilde{M} is glued from infinitely many copies of the \widetilde{M}_i 's with “slabs” $\mathbb{R}^2 \times I$ interpolating between them. Each slab admits a full \mathbb{R}^2 of isometric translations. The pieces \widetilde{M}_i will be called “pieces” and each boundary component of a piece will be called a “flat” and will be oriented as part of the boundary of the piece it belongs to.

Let M' be another such manifold and \widetilde{M}' its universal cover, metrized as above. Kapovich and Leeb [6] show that any quasi-isometry $f: \widetilde{M} \rightarrow \widetilde{M}'$ is a bounded distance from a quasi-isometry that maps slabs to slabs and geometric pieces to geometric pieces. Then a theorem of Schwartz [14] says that f is a uniformly bounded distance from an isometry on each piece \widetilde{M}_i . Our uniform choice of neuterings assures that we can change f by a bounded amount to be an isometry on each piece and an isometry followed by a shear map on each slab (a “shear map” $\mathbb{R}^2 \times I \rightarrow \mathbb{R}^2 \times I$ will mean one of the form $(x, t) \mapsto (x + \rho(t)v, t)$ with $v \in \mathbb{R}^2$ and $\rho: I \rightarrow I$ as described earlier). We then say f is *straightened*.

Consider now the group³

$$\mathcal{I}(\widetilde{M}) := \{f: \widetilde{M} \rightarrow \widetilde{M} \mid f \text{ is a straightened quasi-isometry}\}.$$

$\mathcal{I}(\widetilde{M})$ acts on the set of pieces of \widetilde{M} ; the pieces \widetilde{M}_i that cover a given piece M_i of M are all in one orbit of this action, but the orbit may be larger. The subgroup $\mathcal{I}_{\widetilde{M}_i}(\widetilde{M})$ of $\mathcal{I}(\widetilde{M})$ that stabilizes a fixed piece \widetilde{M}_i of \widetilde{M} acts discretely on \widetilde{M}_i , so $\widetilde{M}_i/\mathcal{I}_{\widetilde{M}_i}(\widetilde{M})$ is an orbifold (which clearly is covered by M_i)

The subgroup $\mathcal{I}_S(\widetilde{M}) \subset \mathcal{I}(\widetilde{M})$ that stabilizes a slab $S \subset \widetilde{M}$ acts on S by isometries composed with shear maps. It acts discretely on each boundary component of S but certainly not on S . But an element that is a finite order rotation on one boundary component must be a similar rotation on the other boundary component (and is in fact finite order on the slab; the group $\mathcal{I}_S(\widetilde{M})$ is abstractly an extension of $\mathbb{Z}^2 \times \mathbb{Z}^2$ by a finite cyclic group F_S of order 1, 2, 3, 4, or 6; the two \mathbb{Z}^2 's are the translation groups for the two boundaries of S).

We form an \mathcal{N} -graph $\Omega(\widetilde{M})$ as follows. The vertices of $\Omega(\widetilde{M})$ correspond to $\mathcal{I}(\widetilde{M})$ -orbits of pieces and the edges correspond to orbits of slabs—the edge determined by a slab connects the vertices determined by the abutting pieces. We label each vertex by the corresponding orbifold $\widetilde{M}_i/\mathcal{I}_{\widetilde{M}_i}(\widetilde{M})$ and each edge by the derivative

³Note that $\mathcal{I}(\widetilde{M}) \rightarrow \mathcal{QI}(\widetilde{M})$ is an isomorphism, where \mathcal{QI} denotes the group of quasi-isometries (defined by identifying maps that differ by a bounded distance).

of the affine map between the flats that bound a corresponding slab S (this map is determined up to the cyclic group F_S).

Proposition 4.1. *The \mathcal{N} -graph $\Omega(\widetilde{M})$ constructed above is determined up to isomorphism by \widetilde{M} . The manifold \widetilde{M} is determined up to bilipschitz diffeomorphism by $\Omega(\widetilde{M})$.*

Proof. The first sentence of the proposition is true by construction.

We describe how to reconstruct \widetilde{M} from $\Omega(\widetilde{M})$. $\Omega(\widetilde{M})$ gives specifications for inductively gluing together pieces that are the universal covers of the orbifolds corresponding to its vertices, and slabs between these pieces, according to a tree: we start with a piece \widetilde{N} corresponding to a vertex v of Γ and glue slabs and adjacent pieces on all boundary components as specified by the outgoing edges at v in $\Omega(\widetilde{M})$, and repeat this process for each adjacent piece, and continue inductively. (There is an underlying tree for this construction which is the universal cover of the graph obtained by replacing each edge of Γ by a countable infinity of edges.) The construction involves choices, since each gluing map is only determined up to a group of isometries of the form $\mathbb{R}^2 \times F_S$, where S is the slab. We need to show that the resulting manifold is well defined up to bilipschitz diffeomorphism.

The constructed manifold X and the original \widetilde{M} can both be constructed in the same way from $\Omega(\widetilde{M})$, but they potentially differ in the choices just mentioned. We can construct a bilipschitz diffeomorphism $f: X \rightarrow \widetilde{M}$ inductively, starting with an isometry from one piece of X to one piece of \widetilde{M} and extending repeatedly over adjacent pieces. At any point in the induction, when extending to an adjacent piece across an adjacent slab S , we use an isometry of the adjacent piece X_i of X to the adjacent piece M_i of \widetilde{M} that takes the boundary component $X_i \cap S$ of X_i to the boundary component $M_i \cap S$ of M_i . Restricted to this boundary component E , this isometry is well defined up to the action of a lattice \mathbb{Z}^2 , so there is a choice that can be extended across the slab S with an amount of shear bounded by the diameter of the torus E/\mathbb{Z}^2 . Since only finitely many isometry classes of such tori occur in the construction, we can inductively construct the desired diffeomorphism using a uniformly bounded amount of shear on slabs. This diffeomorphism therefore has a uniformly bounded bilipschitz constant, as desired. \square

We now describe how the above arguments must be modified if M has arithmetic pieces. Let M_1 be an arithmetic hyperbolic piece which is adjacent to a non-arithmetic hyperbolic piece M_2 . We have already discussed how M_2 is neutered; we take an arbitrary neutering of M_1 (and any other arithmetic pieces) which we will adjust later. As before, the notation M_i refers to the neutered pieces and we glue M from these pieces mediating with toral annuli $T^2 \times I$ between them.

We aim to show that we can adjust the neutering of the arithmetic pieces so that any quasi-isometry of \widetilde{M} can be straightened as in the beginning of this section to be an isometry on pieces and a shear map on slabs.

Consider \widetilde{M}_1 as a subset of \mathbb{H}^3 obtained by removing interiors of infinitely many disjoint horoballs. Schwartz [14] shows that any quasi-isometry of \widetilde{M}_1 is a bounded distance from an isometry of \widetilde{M}_1 to a manifold obtained by changing the sizes of the removed horoballs by a uniformly bounded amount.

In the universal cover \widetilde{M} choose lifts \widetilde{M}_1 and \widetilde{M}_2 glued to the two sides $\partial_1 S$ and $\partial_2 S$ of a slab $S \cong \mathbb{R}^2 \times I$. Consider a quasi-isometry f of \widetilde{M} which maps S to a

bounded Hausdorff distance from itself. We can assume that f is an isometry on \widetilde{M}_2 . By the previous remarks, the map f restricted to \widetilde{M}_1 is a bounded distance from an isometry of \widetilde{M}_1 which moves its boundary component $\partial_1 S$ to a parallel horosphere (if we consider \widetilde{M}_1 as a subset of \mathbb{H}^3); by inverting f if necessary, we can assume the diameter of the horosphere has not decreased. Thus \widetilde{M}_1 can be positioned in $\mathbb{H}^3 = \{(z, y) \in \mathbb{C} \times \mathbb{R} : y > 0\}$ so that $\partial_1 S$ is the horosphere $y = 1$ and $f(\partial_1 S)$ is the horosphere $y = \lambda$ for some $\lambda \leq 1$. Using a smooth isotopy of f which is supported in an ϵ -neighborhood of the region between $f(\partial_1 S)$ and $\partial_1 S$, and which moves $f(\partial_1 S)$ to $\partial_1 S$, we can adjust f to map $\partial_1 S$ to itself. This moves each point of $f(\partial_1 S)$ to its closest point on $\partial_1 S$ by a euclidean similarity, scaling distance uniformly by a factor of λ . The resulting adjusted f is still a quasi-isometry, so restricted to S we then have a quasi-isometry which scales metric on $\partial_1 S$ by λ and is an isometry on $\partial_2 S$. This is only possible if $\lambda = 1$, so f , once straightened on \widetilde{M}_1 , maps $\partial_1 S$ to itself.

Now consider the subgroup of the group of quasi-isometries of \widetilde{M} which takes \widetilde{M}_1 to itself, and just consider its restriction to \widetilde{M}_1 , which we can think of as embedded in \mathbb{H}^3 . By straightening, we have a group of isometries of \mathbb{H}^3 which preserves a family of disjoint horoballs (the ones that are bounded by images of $\partial_1 S$). Any subgroup of $\text{Isom}(\mathbb{H}^3)$ which preserves an infinite family of disjoint horoballs is discrete. Thus, from the point of view of the construction above, M_1 behaves like a non-arithmetic piece. By repeating the argument, this behavior propagates to any adjacent arithmetic pieces, hence, so long as at least one piece is non-arithmetic, the construction of the graph $\Omega(\widetilde{M})$ goes through as before and the proof of Proposition 4.1 extends.

We have not yet proved Theorem 3.7, which says that there is a unique minimal graph in each bisimilarity class. The proof that $\Omega(\widetilde{M})$ is minimal will follow.

Proof of Theorem 3.7. The construction of the proof of Proposition 4.1 works for any \mathcal{N} -graph Γ , gluing together infinitely many copies of the universal covers \widetilde{N}_i of the orbifolds that label the vertices, with slabs between them, according to an infinite tree (the universal cover of the graph obtained by replacing each edge of Γ by countable-infinitely many). We get a simply connected Riemannian manifold $X(\Gamma)$ which, as long as Γ is finite, is well defined up to bilipschitz diffeomorphism by the same argument as before.

The group of straightened self-diffeomorphisms $\mathcal{I}(X(\Gamma))$, when restricted to a piece \widetilde{N}_i , includes the covering transformations for the covering $\widetilde{N}_i \rightarrow N_i$. It follows that the construction of a \mathcal{N} -graph from $X(\Gamma)$, as given in the first part of this section, yields an \mathcal{N} -graph $\Omega(X(\Gamma))$ together with a morphism $\Gamma \rightarrow \Omega(X(\Gamma))$.

If $\Gamma \rightarrow \Gamma'$ is a morphism of \mathcal{N} -graphs, then $X(\Gamma)$ and $X(\Gamma')$ are bilipschitz diffeomorphic, since the instructions for assembling them are equivalent. Hence $\Omega(X(\Gamma)) = \Omega(X(\Gamma'))$; call this graph $m(\Gamma)$. Since $m(\Gamma) = m(\Gamma')$ and the existence of a morphism generates the relation of bisimilarity of \mathcal{N} -graphs, $m(\Gamma)$ is the same for every graph in the bisimilarity class. It is also the target of a morphism from every Γ in this class. Thus it is the unique minimal element in the class. Thus Theorem 3.7 is proved. \square

Proposition 4.2. *For any \mathcal{N} -manifold M , $\Omega(\widetilde{M})$ is balanced, and is the minimal \mathcal{N} -graph in the bisimilarity class of the \mathcal{N} -graph $\Gamma(M)$ associated with M (Definition 3.3).*

Proof. By construction, there is a morphism $\Gamma(M) \rightarrow \Omega(\widetilde{M})$. In particular, $\Omega(\widetilde{M})$ is in the bisimilarity class of $\Gamma(M)$. It is balanced by Theorem 3.8, since $\Gamma(M)$ is integral.

If M is an \mathcal{N} -orbifold and $\Gamma = \Gamma(M)$ its associated integral \mathcal{N} -graph, then $X(\Gamma)$ is bilipschitz homeomorphic to \widetilde{M} , so $\Omega(\widetilde{M}) = \Omega(X(\Gamma)) = m(\Gamma)$, and is hence minimal by the previous proof, so Proposition 4.2 now follows. \square

Proof of Classification Theorem for \mathcal{N} -manifolds. Since \widetilde{M} is quasi-isometric to $\pi_1(M)$, a quasi-isometry between fundamental groups of M and M' induces a quasi-isometry $\widetilde{M} \rightarrow \widetilde{M}'$. We have already explained how this can then be straightened and thus give an isomorphism of the corresponding minimal \mathcal{N} -graphs. \square

For the Classification Theorem to be a complete classification of quasi-isometry types of fundamental groups of good manifolds we will need to know that every minimal balanced \mathcal{N} -graph is realized by an \mathcal{N} -manifold. This is the content of the Realization Theorem which we prove later.

Remark 4.3. In the above proof of Theorem 3.7, we showed that \mathcal{N} -manifolds M, M' with bisimilar \mathcal{N} -graphs Γ, Γ' have bilipschitz diffeomorphic universal covers, since the common $X(\Gamma) = X(\Gamma')$ is a model for the universal covers. A slight modification of the proof shows that it is also true that \mathcal{H} -manifolds with bisimilar graphs have bilipschitz universal covers. For an \mathcal{H} -manifold which is not an \mathcal{N} -manifold there is no longer a unique minimal \mathcal{H} -graph in the bisimilarity class. The existence of this unique minimum was essential to the above proof of the converse: that bilipschitz universal covers of \mathcal{N} -manifolds implies bisimilar graphs.

In fact, we prove in [2] that in the “all-arithmetic” case every \mathcal{H} -graph has infinitely many minimal elements in its bisimilarity class.

5. COVERS AND RFCH

For our realization theorem, and also for our commensurability theorem, we need to produce covers of 3-manifolds with prescribed boundary behavior. The covers we need can be summarized by the following purely topological conjecture, which we find is of independent interest.

Cusp Covering Conjecture 5.1 (CCC_n). *Let M be a hyperbolic n -manifold. Then there exists an assignment of a sublattice Λ_C of $\pi_1(C)$ for each cusp C of M such that, for any assignment of a sublattice $\Lambda'_C \subset \Lambda_C$ for each C , there exists a finite cover M' of M whose cusps covering each cusp C of M are the covers determined by Λ'_C .*

We only need this conjecture for $n = 3$, which, as we mention in the Introduction, follows from work of Dani Wise [15, Corollary 18.11].

We now show that the general version of the conjecture follows from the well-known residual finiteness conjecture for hyperbolic groups (RFCH).

Theorem 5.2. *The residual finiteness conjecture for hyperbolic groups (RFCH) implies CCC_n for all n .*

Proof. The Dehn surgery theorem for relatively hyperbolic groups of Osin [10] and Groves and Manning [3] guarantees the existence of sublattices Λ_C of $\pi_1(C)$ for each cusp so that, given any subgroups $\Lambda'_C \subset \Lambda_C$ for each C , the result of adding relations to $\pi_1(M)$ which kill each Λ'_C gives a group G into which the groups

$\pi_1(C)/\Lambda'_C$ inject and which is relatively hyperbolic relative to these subgroups. If the Λ'_C are sub-lattices, then G is relatively hyperbolic relative to finite (hence hyperbolic) subgroups, and is hence itself hyperbolic. By RFCH we may assume G is residually finite, so there is a homomorphism of G to a finite group H such that each of the finite subgroups $\pi_1(C)/\Lambda'_C$ of G injects. The kernel K of the composite homomorphism $\pi_1(M) \rightarrow G \rightarrow H$ thus intersects each $\pi_1(C)$ in the subgroup Λ'_C . The covering of M determined by K therefore has the desired property. \square

Let Γ be a \mathcal{N} -graph. For an edge e of Γ let T_e be the tangent space of the cusp orbifold corresponding to the start of e . We can identify T_e with $T_{\bar{e}}$ using the linear map ℓ_e , so we will generally not distinguish T_e and $T_{\bar{e}}$. Then T_e contains two \mathbb{Z} -lattices, the underlying lattices for the cusp orbifolds associated with the initial and terminal segments of e , and we will denote their intersection by Λ_e . Thus, a torus T_e/Λ is a common cover of the cusp orbifolds at the two ends of e if and only if $\Lambda \subset \Lambda_e$.

From CCC₃, we can now choose a sublattice Λ'_e of Λ_e for each edge e of Γ such that for each vertex v of Γ the corresponding orbifold N_v has a cover M_v with the following property:

Property 5.3. For each cusp of N_v corresponding to an edge e departing v , all cusps of M_v which cover it are of type T_e/Λ'_e .

This property is precisely the consequence of the CCC₃ that we use in our proofs of the Realization and Commensurability Theorems.

6. REALIZING GRAPHS

We already proved the first sentence of the Realization Theorem of the Introduction to this paper. We now prove the second part.

Proof of Realizability Theorem for \mathcal{H} -manifolds. Let Γ be a balanced \mathcal{H} -graph. We want to show there is some \mathcal{H} -manifold which realizes a graph in its bisimilarity class. As pointed out in the previous section, this is equivalent to finding an integral \mathcal{H} -graph in the bisimilarity class.

Using CCC₃, after we choose a sublattice Λ'_e of Λ_e for each edge e of Γ , we may assume Property 5.3 holds.

Let d_v be the degree of the cover $M_v \rightarrow N_v$. For an edge e departing v let d_e be the degree of the corresponding cover of cusp orbifolds of M_v and N_v , i.e., the index of the lattice Λ'_e in the fundamental group of the boundary component C_e of N_v corresponding to e (this is slightly different from the usage in the proof of Theorem 3.8). Since the cusps of the M_v corresponding to the two ends of e are equal, we have

$$(5) \quad d_e \delta_e = d_{\bar{e}}.$$

(Recall that δ_e denotes the determinant of the linear map ℓ_e .)

Since Γ is balanced, we can assign a positive rational number $m(v)$ to each vertex with the property that for any edge e one has $m(\tau e) = \delta_e m(\iota e)$. Thus, by (5),

$$(6) \quad \frac{m(\tau e)}{d_{\bar{e}}} = \frac{m(\iota e)}{d_e}.$$

Choose a positive integer b such that $n(v) := \frac{bm(v)}{d_v}$ is integral for every vertex of Γ (so b is some multiple of the lcm of the denominators of the numbers $\frac{m(v)}{d_v}$). Let

M'_v be the disjoint union of $n(v)$ copies of M_v , so M'_v is a $bm(v)$ -fold cover of N_v . Let $\pi_v: M'_v \rightarrow N_v$ be the covering map.

For an edge e of Γ from $v = \iota e$ to $w = \tau e$, the number of boundary components of M'_v covering the boundary component C_e of N_v corresponding to e is $bm(v)/d_e$. By (6) this equals $bm(w)/d_{\bar{e}}$, which is the number of boundary components of M'_w covering the boundary component $C_{\bar{e}}$ of N_w . Thus $\pi_v^{-1}C_e$ and $\pi_w^{-1}C_{\bar{e}}$ have the same number of components, and each component is T_e/Λ' , so we can glue M'_v to M'_w along these boundary components using any one-one matching between them. Doing this for every edge gives a manifold M whose \mathcal{H} -graph has a morphism to Γ ; if M is disconnected, replace it by a component (but one can always do the construction so that M is connected). This proves the theorem. \square

7. COMMENSURABILITY

Proof of Commensurability Theorem. Let M_1 and M_2 be two \mathcal{N} -manifolds whose \mathcal{N} -graphs are bisimilar. Assume that their common minimal \mathcal{N} -graph is a tree and all the vertex labels are manifolds. We want to show that M_1 and M_2 are commensurable.

Let Γ be the common minimal \mathcal{N} -graph for M_1 and M_2 . Using CCC_3 we may assume Property 5.3 holds. Accordingly, for each vertex of Γ we may take a common cover N'_v of all the pieces of M_1 that cover N_v and then choose the lattices Λ'_e , as in the previous section, to be subgroups of the cusp groups of these manifolds N'_v . In this way we arrange that the pieces M_v of the manifold M constructed in that section are covers of N'_v , and hence of the pieces of M_1 . Elementary arithmetic as in the proof of the Realization Theorem shows that there is a b_0 so that if the number b of that proof is a multiple of b_0 then we can choose the gluing in that proof to make M a covering space of M_1 . Call the resulting manifold M'_1 . If we initially choose N'_v to also cover the type v pieces of M_2 then we can also construct a covering space M'_2 of M_2 out of copies of pieces M_v . The decompositions of M'_1 and M'_2 then give \mathcal{N} -graphs whose vertices are labeled by M_v 's and whose edges are labeled by \mathbb{Z} -isomorphisms. It suffices to show that M'_1 and M'_2 are commensurable.

We can encode the information needed to construct M'_1 in a simplified version $\Gamma_0(M'_1)$ of its \mathcal{N} -graph. The underlying graph is still just the graph describing the decomposition of M'_1 into pieces, but the labeling is simplified as follows. Each vertex w of $\Gamma_0(M'_1)$ corresponds to a copy of some M_v , which is a normal covering of the orbifold N_v . We say vertex w is of *type* v . The covering transformation group for the covering $M_v \rightarrow N_v$ induces a permutation group P_v on the edges of $\Gamma_0(M'_1)$ exiting w . We record the type and permutation action for each vertex of $\Gamma_0(M'_1)$. It is easy to see that the graph $\Gamma_0(M'_1)$ with these data records enough information to reconstruct M'_1 up to diffeomorphism from its pieces. A covering of such a graph induces a covering of the same degree for the manifolds they encode.

A graph with fixed permutation actions at vertices as above is called *symmetry restricted* in [8]. The graphs $\Gamma_0(M'_1)$ and $\Gamma_0(M'_2)$ have the same universal covering as symmetry restricted graphs.

So we would like to know that when two such graphs have a common universal covering, then they have a common finite covering. A generalization, proved in [8] of Leighton's theorem [7] (which deals with graphs without the extra structure) shows that this is true if the underlying graph is a tree, so we are done. \square

The results of [8] allow one to carry out the above proof under slightly weaker assumptions on the minimal graph than being a tree, but not sufficiently general to warrant going into details.

8. ADDING SEIFERT FIBERED PIECES

We will modify Definition 3.1 to allow the inclusion of Seifert fibered space pieces. The graphs we use are called \mathcal{NS} -graphs, and we define them below.

We will need to consider Seifert fibered orbifolds among the pieces, so we first describe a coarse classification into types. As always, the manifolds and orbifolds we consider are oriented. We will distinguish two types: the oriented Seifert fibered orbifold N is type “o” or “n” according as the Seifert fibers can be consistently oriented or not. N is type “o” if and only if the base orbifold S of the Seifert fibration is orientable. This can fail in two ways: the topological surface underlying S may be non-orientable, or S may be non-orientable because it has mirrors. The latter arises when parts of N look locally like a Seifert fibered solid torus $D^2 \times S^1$ factored by the involution $(z_1, z_2) \mapsto (\bar{z}_1, \bar{z}_2)$ (using coordinates in \mathbb{C}^2 with $|z_1| \leq 1$ and $|z_2| = 1$). The fibers with $z_1 \in \mathbb{R}$ in this local description are intervals (orbifolds of the form $S^1/(\mathbb{Z}/2)$), and the set of base points of such interval fibers of N form mirror curves in S which are intervals and/or circles embedded in the topological boundary of S . If any of these mirror curves are intervals, so they merge with part of the true boundary of S (image of boundary of N), then the corresponding boundary component of N is a pillow orbifold (topologically a 2-sphere, with four 2-orbifold points).

If two type “o” Seifert fibered pieces in the decomposition of M are adjacent along a torus, and we have oriented their Seifert fibers, then there is a sense in which these orientations are compatible or not. We say the orientation is *compatible* or *positive* if, when viewing the torus from one side, the intersection number in the torus of a fiber from the near side with a fiber from the far side is positive. Note that this is well defined, since if we view the torus from the other side, both its orientation and the order of the two curves being intersected have changed, so the intersection number is unchanged.

We define \mathcal{NS} -graphs below, with the geometric meanings of the new ingredients in square brackets. But there is a caveat to these descriptions. Just as an \mathcal{N} -graph can be associated to the geometric decomposition of an \mathcal{N} -manifold, an \mathcal{NS} -graph can be associated with the decomposition of a good manifold or orbifold M (see the Introduction for the definition of “good”). However, the \mathcal{N} -graph associated with a geometric decomposition is a special kind of \mathcal{N} -graph (it is “integral” in the terminology of Section 3), and an \mathcal{NS} -graph coming from the geometric decomposition of a good manifold is similarly special. The geometric explanations therefore only match precisely for special cases of \mathcal{NS} -graphs. Recall that a subgraph Γ' of Γ is full if for each vertices $a, b \in \Gamma'$ any edges connecting them in Γ , are also in Γ' .

Definition 8.1. An \mathcal{NS} -graph Γ is a finite connected graph with decorations on its vertices and edges as follows:

- (1) Vertices are partitioned into two types: *hyperbolic vertices* and *Seifert vertices*. The full subgraphs of Γ determined by the hyperbolic vertices, respectively the Seifert vertices, are called the *hyperbolic subgraph*, respectively the *Seifert subgraph*.

- (2) The hyperbolic subgraph is labeled as in Definition 3.1, so that each of its components is an \mathcal{N} -graph. In particular, each hyperbolic vertex v is labeled by a hyperbolic orbifold N_v and there is a map $e \mapsto C_e$ from the set of directed edges e exiting that vertex to the set of cusp orbifolds C_e of N_v . This map is defined on the set of *all* edges exiting e , not just the edges in the hyperbolic subgraph. As before, it is injective, except that an edge which begins and ends at the same vertex may have $C_e = C_{\bar{e}}$.
- (3) Each Seifert vertex is labeled by one of two colors, black or white [for an \mathcal{NS} -graph coming from a geometric decomposition this encodes whether the Seifert fibered piece in the geometric decomposition of M contains boundary components of the ambient 3-manifold or not, as in [1]]. It is also labeled by a Seifert fibration type “o” or “n”, as described above.
- (4) For an edge e starting at a Seifert vertex the group $F_e = F_{\bar{e}}$ is $\{1\}$ or $\{\pm 1\}$. If F_e is $\{\pm 1\}$ then the Seifert vertex is type “n”.
- (5) For each edge e from a hyperbolic vertex to a Seifert vertex the group F_e associated to the cusp section C_e is either trivial or $\{\pm 1\}$. The edge is labeled by a non-zero rational vector in TC_e called the *slope* s_e , determined up to the action of F_e . [s_e encodes the direction and length of the fibers of the adjacent Seifert fibered piece.]
- (6) Each edge e connecting a type “o” Seifert vertex with a type “o” Seifert vertex or a hyperbolic vertex has a *sign label* $\epsilon_e = \pm 1$ with $\epsilon_e = \epsilon_{\bar{e}}$ [this describes compatibility of orientations of Seifert fibers of adjacent pieces or—if the edge connects a Seifert and a hyperbolic vertex—of Seifert fiber and slope].
- (7) The data described in items (5) and (6) are subject to the equivalence relation generated by the following moves:
 - (a) For any type “o” Seifert vertex the signs at all edges adjacent to it may be multiplied by -1 [reversal of orientation of the Seifert fibers].
 - (b) The slope s_e of item (5) can be multiplied by -1 while simultaneously multiplying ϵ_e by -1 .
 - (c) For any Seifert vertex, the slopes at all adjacent hyperbolic vertices may be multiplied by a fixed non-zero rational number.

Note that the data encoded by the sign weights modulo the equivalence relation of item (7a) are equivalent to an element of $H^1(\Gamma \setminus \Gamma_n, \Gamma_h; \mathbb{Z}/2)$, where Γ_h is the hyperbolic subgraph and Γ_n the full subgraph on Seifert vertices of type “n”.

We define a *morphism of \mathcal{NS} -graphs*, $\pi: \Gamma \rightarrow \Gamma'$, to be an open graph homomorphism which restricts to a \mathcal{N} -graph morphism of the hyperbolic subgraphs (Definition 3.4), preserves the black/white coloring on the Seifert vertices, and on the edges between hyperbolic and Seifert vertices preserves the slope (in the sense that the slope at the image edge is the image under the tangent map on cusp orbifolds of the slope at the source edge). Moreover, it must map type “n” vertices to type “n” vertices, and when an “o” vertex v is mapped to an “n” vertex w , then the preimage of each edge at w must either include edges of different signs, or an edge terminating in a type “n” Seifert vertex.

As in Section 3, the existence of morphisms between \mathcal{NS} -graphs generates an equivalence relation which we call *bisimilarity*.

Proof of Classification Theorem. To prove the Classification Theorem we will show that each bisimilarity class of \mathcal{NS} -graphs has a minimal element, and if a \mathcal{NS} -graph

comes from a non-geometric manifold M then the minimal \mathcal{NS} -graph determines and is determined by \widetilde{M} up to quasi-isometry.

We first explain why the qi-type of the universal cover \widetilde{M} (or equivalently of $\pi_1(M)$) determines a minimal \mathcal{NS} -graph $\Omega(\widetilde{M})$.

As in Section 4, we can straighten any quasi-isometry $\widetilde{M} \rightarrow \widetilde{M}'$ and assume it takes geometric pieces to geometric pieces and slabs to slabs. We may also assume it is an isometry on hyperbolic pieces. A Seifert piece in \widetilde{M} is bi-Lipschitz homeomorphic to a fattened tree times \mathbb{R} , so the fibration by \mathbb{R} fibers is coarsely preserved, and we can straighten it so it is actually preserved. Moreover, if an adjacent piece is hyperbolic, then the straightened quasi-isometry is an isometry on the corresponding flat, so the affine structure on fibers of the Seifert piece is coarsely preserved, and we can straighten so that it is actually preserved. However, where a Seifert piece is adjacent to a Seifert piece the $\mathbb{R} \times \mathbb{R}$ product structure on the corresponding flat (given by Seifert fibers on the two sides) is coarsely preserved, but the affine structures on the \mathbb{R} fibers need only be preserved up to quasi-isometry.

Considering straightened quasi-isometries in the above sense, we denote again

$$\mathcal{I}(\widetilde{M}) := \{f: \widetilde{M} \rightarrow \widetilde{M} \mid f \text{ is a straightened quasi-isometry}\}.$$

As in Section 4, the underlying graph for our minimal \mathcal{NS} -graph $\Omega(\widetilde{M})$ has a vertex for each orbit of the action of $\mathcal{I}(\widetilde{M})$ on the set of pieces of \widetilde{M} and edge for each orbit of the action on the set of slabs. The labeling of the hyperbolic subgraph is as before; in particular, any vertex corresponding to an orbit of hyperbolic pieces is labeled by the hyperbolic orbifold obtained by quotienting a representative piece in the orbit by its isotropy subgroup in $\mathcal{I}(\widetilde{M})$. A Seifert vertex of the \mathcal{NS} -graph is of type “n” if some element of $\mathcal{I}(\widetilde{M})$ takes a corresponding Seifert fibered piece to itself reversing orientations of fibers, and is otherwise of type “o”. For each type “o” vertex we choose an orientation of the fibers of one piece in the corresponding orbit and then extend equivariantly to the other pieces in the orbit.

For a Seifert vertex adjacent to at least one hyperbolic vertex the fibers of the corresponding pieces in \widetilde{M} carry an affine structure which is defined up to affine scaling. We choose a specific scale for each such vertex, so we can speak of length along fibers, and then the *slope* of item (5) of Definition 8.1 is given by a tangent vector of unit length, viewed in the adjacent cusp.

The sign weights of item (6) of Definition 8.1 are then defined, and item (7) of that definition reflects the choices of orientation and scale which were made.

By construction, the isomorphism type of $\Omega(\widetilde{M})$ is determined by \widetilde{M} and thus two manifolds with quasi-isometric fundamental group have the same associated graph. It remains to show that the isomorphism type of $\Omega(\widetilde{M})$ determines the bilipschitz homeomorphism type of \widetilde{M} .

Construct a labeled graph $\widetilde{\Omega}$ from $\Omega = \Omega(\widetilde{M})$ by first replacing each edge of Ω by infinitely many edges, keeping the weights on edges, but adding sign weights to edges at type “n” vertices with infinitely many of each sign, and then taking the universal cover of the resulting weighted graph. Finally, “o” and “n” labels are now irrelevant and can be removed.

To associate a manifold $X = X(\widetilde{\Omega})$ to this labeled graph, we must glue together appropriate pieces according to the tree $\widetilde{\Omega}$, with appropriate choices for the gluing between slabs. The pieces for hyperbolic vertices will be universal covers of the

hyperbolic orbifolds which label them, while for a Seifert vertex we take the universal cover of some fixed Seifert fibered manifold with base of hyperbolic type and having a boundary component for each incident edge in Ω and—if the vertex is a black vertex—an additional boundary component to contribute to boundary of X . (The universal cover Y of this Seifert piece is then a fattened tree times \mathbb{R} , and, as in [1], it is in fact only important that there is a bound B such that for each boundary component of Y there are boundary components of all “types” within distance B of the give boundary component.)

The choices in gluing depend on the types of the abutting pieces: Between hyperbolic pieces, the gluing map is, as before, determined up to a group of isometries of the form $\mathbb{R}^2 \times F_S$, where S is the slab. Between Seifert fibered pieces the gluing will be an affine map such that the fibers from the two pieces then intersect in the intervening flat with sign given by the sign label of the edge. Finally, between a Seifert fibered and a hyperbolic piece the gluing will be an affine map matching unit tangent vector along fibers with the slope vector for the hyperbolic piece. For each edge of Ω we make a fixed choice of how to do the gluing subject to the above constraints and do it this way for every corresponding edge of $\tilde{\Omega}$.

To complete the proof, it remains to show that, independent of these choices, there exists a bilipschitz homeomorphism from \tilde{M} to X .

As in the proof of Proposition 4.1, the desired bilipschitz homeomorphism is built inductively, starting with a homeomorphism from one piece of \tilde{M} to a piece of X and then extending via adjacent slabs to adjacent pieces. There are four cases: (1), when both adjacent pieces are hyperbolic, this is exactly the case of Proposition 4.1; (2), when both pieces are Seifert fibered; (3), extending from a hyperbolic piece to an adjacent Seifert fibered piece; and (4), extending from a Seifert fibered piece to an adjacent hyperbolic piece. For Case (2) we use [1, Theorem 1.3] (as in the proof of [1, Theorem 3.2]) to extend over the adjacent Seifert fibered piece, respecting the “types” of boundary components (i.e., belonging to boundary of M or not, and if not, then the “type” is given by the edge of Ω that the boundary component corresponds to). Case (3) is essentially the same argument, and Case (4) is immediate.

Thus we obtain the desired bilipschitz homeomorphism, completing the proof. \square

If we remove the restriction that the components of the hyperbolic subgraph be \mathcal{N} -graphs in the Definition 8.1 of an \mathcal{NS} -graph we get the definition of an \mathcal{HS} -graph. Any irreducible non-geometric 3-manifold has an associated \mathcal{HS} -graph.

Proof of Realization Theorem. The construction of $\Omega(\tilde{M})$, given above, has built in a morphism from the \mathcal{HS} -graph associated to M . Since the graph associated to M is balanced, it follows from Proposition 4.2 that $\Omega(\tilde{M})$ is balanced.

The balanced condition is only a constraint on \mathcal{H} -graph components of an \mathcal{HS} -graph. Thus, from the case of \mathcal{H} -graphs which we established in Section 6, we may conclude that every balanced \mathcal{HS} -graph is in the bisimilarity class of some \mathcal{HS} -graph of an \mathcal{HS} -manifold. In particular, a minimal balanced \mathcal{NS} -graph is the minimal graph for some quasi-isometry class of good 3-manifold groups. \square

REFERENCES

- [1] J. Behrstock and W.D. Neumann, Quasi-isometric classification of graph manifolds groups. *Duke Math. J.* **141** (2008), 217–240.
- [2] J. Behrstock and W.D. Neumann, in preparation.

- [3] D. Groves and J.F. Manning, Dehn filling in relatively hyperbolic groups. *Israel J. Math.* **168** (2008), 317–429.
- [4] Emily Hamilton, Abelian subgroup separability of Haken 3-manifolds and closed hyperbolic n -orbifolds. *Proc. London Math. Soc.* **83** (2001), 626–646.
- [5] J. Hempel, Residual finiteness for 3-manifolds, in *Combinatorial group theory and topology (Alta, Utah, 1984)*, *Ann. of Math. Stud.*, **111** (Princeton Univ. Press, Princeton, NJ, 1987) 379–396
- [6] M. Kapovich and B. Leeb. Quasi-isometries preserve the geometric decomposition of Haken manifolds. *Invent. Math.* **128** (1997), 393–416.
- [7] Frank Thomson Leighton, Finite Common Coverings of Graphs. *J. Comb. Theory, Series B* **33** (1982), 231–238.
- [8] Walter D Neumann, On Leighton’s graph covering theorem. *Groups, Geometry, and Dynamics* **4** (2010), 863–872.
- [9] Walter D. Neumann. Commensurability and virtual fibration for graph manifolds. *Topology* **36** (1997), 355–378.
- [10] Denis V. Osin, Peripheral fillings of relatively hyperbolic groups, *Invent. Math.* **167** (2007), no. 2, 295–326.
- [11] G. Perelman. The entropy formula for the Ricci flow and its geometric applications. Preprint, [ARXIV:MATH.DG/0211159](https://arxiv.org/abs/math/0211159), 2002.
- [12] G. Perelman. Ricci flow with surgery on three-manifolds. Preprint, [ARXIV:MATH.DG/0303109](https://arxiv.org/abs/math/0303109), 2003.
- [13] G. Perelman. Finite extinction time for the solutions to the Ricci flow on certain three-manifolds. Preprint, [ARXIV:MATH.DG/0307245](https://arxiv.org/abs/math/0307245), 2003.
- [14] R. Schwartz. The quasi-isometry classification of rank one lattices. *IHES Sci. Publ. Math.*, **82** (1996) 133–168.
- [15] Dani Wise. The structure of groups with a quasiconvex hierarchy. Preprint 2010.
- [16] Dani Wise. Research announcement: The structure of groups with a quasiconvex hierarchy. *Electron. Res. Announc. Math. Sci.* **16** (2009), 44–55.

DEPARTMENT OF MATHEMATICS, LEHMAN COLLEGE, CUNY
E-mail address: jason.behrstock@lehman.cuny.edu

DEPARTMENT OF MATHEMATICS, BARNARD COLLEGE, COLUMBIA UNIVERSITY
E-mail address: neumann@math.columbia.edu