ISSN 1364-0380 (on line) 1465-3060 (printed)

Geometry & Topology
Volume 6 (2002) 1–26
Published: 16 January 2002

Algorithmic detection and description of hyperbolic structures on closed 3–manifolds with solvable word problem

Jason Manning

Department of Mathematics, University of California at Santa Barbara
Santa Barbara, CA 93106, USA

Email: `manning@math.ucsb.edu`

Abstract

We outline a rigorous algorithm, first suggested by Casson, for determining whether a closed orientable 3-manifold M is hyperbolic, and to compute the hyperbolic structure, if one exists. The algorithm requires that a procedure has been given to solve the word problem in $\pi_1 M$.

AMS Classification numbers Primary: 57M50

Secondary: 20F10

Keywords: 3–manifold, Kleinian group, word problem, recognition problem, geometric structure

Proposed: David Gabai
Seconded: Jean-Pierre Otal, Robion Kirby

Received: 20 February 2001
Revised: 26 October 2001

1 Introduction

Unless otherwise noted, all manifolds are assumed to be orientable. We begin with a couple of definitions:

Definition 1.1 We say that a 3–manifold M has a *geometric structure* (or that M is *geometric*) if there is a homeomorphism $f\colon \mathrm{Interior}(M) \to N$ where N is a complete homogeneous Riemannian manifold which is locally isometric to one of the eight model geometries, which are discussed in [20]. If N is locally isometric to hyperbolic 3–space, we say M has a *hyperbolic structure*. The terms *spherical structure*, *Euclidean structure*, and so on are similarly defined.

If there is an algorithm to construct N (for instance as a finite number of charts or as a convex polyhedron in the model geometry with face identifications) and an algorithm to construct the homeomorphism $f\colon M \to N$, we say that the geometric structure on M is *algorithmically constructible.*

Remark 1.2 If two triangulated compact 3–manifolds M and N are homeomorphic, then it is shown in [15] (see also [4]) that the given triangulations admit isomorphic subdivisions. For any n, there are finitely many subdivisions of the triangulations of N and M with n tetrahedra, so we can find them all and test each pair for isomorphism. Eventually we must find subdivisions which are isomorphic, and the isomorphism gives a homeomorphism between the two manifolds. Thus to prove that the geometric structure on a triangulated 3–manifold M is algorithmically constructible it is necessary only to produce a triangulation of N.

Definition 1.3 Let p be some statement. An algorithm *decides if p* if whenever p is true, the algorithm reports that p is true, and whenever p is false, the algorithm remains silent. An algorithm *decides whether or not p* if whenever p is true the algorithm reports that p is true, and whenever p is false the algorithm reports that p is false.

It is an open question whether there is an algorithm to decide whether or not a closed triangulated 3–manifold has a hyperbolic structure. In this paper we prove:

Theorem 5.1 *There exists an algorithm which will, given a triangulated orientable closed 3-manifold M and a solution to the word problem in $\pi_1 M$, decide whether or not M has a hyperbolic structure.*

The algorithm described gives us detailed information about the hyperbolic structure, and we are able to show further that:

Theorem 5.2 *If M is a triangulated closed orientable 3–manifold which has a hyperbolic structure, then the hyperbolic structure is algorithmically constructible.*

1.1 Context

Several classes of manifolds are known to have fundamental groups with solvable word problems; Waldhausen solved the word problem for Haken and virtually Haken manifolds [28], Niblo and Reeves for manifolds (of any dimension) with a nonpositively curved cubing [16]. Skinner has solved the word problem in a class of non-Haken manifolds containing immersed surfaces of a particular type [24]. Manifolds with word hyperbolic fundamental group also have solvable word problem [7]. Lackenby [11] has shown that there are many 3-manifolds with word hyperbolic fundamental group, not all of which are known to be hyperbolic. If Thurston's geometrization conjecture [27] is true, then these manifolds are in fact hyperbolic.

If M is any 3–manifold satisfying the geometrization conjecture, a solution to the word problem in its fundamental group may be found. Indeed, the manifold may be first decomposed into prime pieces (see for instance [10]), whose fundamental groups form the free factors of $\pi_1 M$, and each of these pieces either has automatic fundamental group or is modeled on nilgeometry or solvgeometry (see [7], Chapter 12 for a proof). But any compact manifold modeled on nilgeometry or solvgeometry is finitely covered by a circle bundle over a torus or a torus bundle over a circle, respectively (see [20]). In particular, such a manifold is virtually Haken, so Waldhausen's solution may be applied. If the geometry of the individual pieces is unknown at the outset, the search for an automatic structure and the search for a finite cover which is Haken may be conducted in parallel on each piece until one or the other is found. A procedure to find the automatic structure of a group is given in [7]; one to determine whether a 3–manifold is Haken is given in [10].

Although no known algorithm decides whether or not a 3–manifold is geometric, some particular geometric manifolds can be recognized algorithmically. In [26] Thompson gives an algorithm to recognize the 3–sphere (see also [19]). Similar methods provide recognition algorithms for the lens spaces [25]. Given the first homology of M (which is easily computed) there is a finite list of lens spaces which M could be homeomorphic to, so we can determine whether a manifold is a lens space. Rubinstein and Rannard have claimed to be able to recognize small Seifert fibered spaces. Note that the homeomorphism problem for Haken 3-manifolds is decidable by work of Haken and Hemion [29],[9]. The homeomorphism problem is also decidable for "rigid weakly geometric" 3-manifolds by the work of Sela [23]. An irreducible 3–manifold M is *rigid weakly geometric* if one of the following holds: (1) M is Haken, (2) M is geometric, or (3) $\pi_1 M$ is word hyperbolic and any irreducible 3–manifold homotopy equivalent to M is in fact homeomorphic to M.

If a closed orientable 3-manifold is known beforehand to satisfy the geometrization conjecture, then the problem of determining its decomposition and the geometric structures on the pieces is completely solved. This is the case if, for instance, the manifold is Haken (see [27]) or admits a symmetry with ≥ 1-dimensional fixed point set (see [5]). For suppose M is such a manifold, which we may suppose is irreducible. If M is Haken, there is an algorithm (algorithm 8.1 of [10]) to find the characteristic submanifold Σ. There are then three cases: (1) $\Sigma = \emptyset$, (2) $\Sigma = M$ (in which case M is a Seifert fibered space), and (3) Σ is neither empty nor all of M.

In case (1) M must be hyperbolic or a small Seifert fiber space. Sela gives a way to determine whether M is a small Seifert fiber space, and if so which one (details are in section 10 of [23] and in Sela's Ph.D. thesis [22]). If M is hyperbolic, the fundamental group is automatic, so a solution to the word problem may be readily found. The methods of the present paper may then be used to find the hyperbolic structure on M.

In case (2), the algorithm from [10] gives enough information to determine the homeomorphism type of the Seifert fiber space, since it constructs the fibering explicitly. The geometric structure is then immediate.

In case (3), we first check whether $\partial\Sigma$ cuts M into pieces which are all homeomorphic to $T^2 \times I$ or an I-bundle over the Klein bottle. In this case M has a solvgeometric structure which can be easily determined. Otherwise, $M \setminus \partial\Sigma$ consists of pieces which are geometric or I–bundles which separate the geometric pieces from one another. Each geometric piece is either Seifert fibered or hyperbolic. The Seifert fibered pieces can be handled as in case (2), and the hyperbolic structures on the remaining pieces can be computed by a variety of methods, for instance by a variation on the methods of the present paper.

Alternatively, results of and Petronio and Weeks [17] have shown that a manifold with torus boundary components admits a finite volume hyperbolic structure if and only if Thurston's hyperbolization equations for some ideal triangulation have a solution corresponding to a combination of positively oriented and flat tetrahedra. Any ideal triangulation is accessible from any other by a sequence of finitely many moves, so this gives an algorithm to find the hyperbolic structure, if one is known to exist.

We summarize the preceding discussion as:

Theorem 1.4 *If a closed orientable 3–manifold is known to satisfy the geometrization conjecture, there are algorithms to find its canonical decomposition, and the geometric structures on the pieces are algorithmically constructible.*

1.2 Outline

We describe an algorithm to decide if a closed triangulated orientable 3–manifold M is hyperbolic, assuming that the word problem in $\pi_1 M$ can be solved. The key to the effectiveness of the algorithm is the following theorem of Gabai, Meyerhoff and N Thurston.

Theorem 1.5 [8] *Let $f\colon M \to N$ be a homotopy equivalence between closed 3–manifolds, where M is irreducible and N is hyperbolic. Then f is homotopic to a homeomorphism.*

This theorem implies that in order to determine that a closed 3-manifold M is hyperbolic, we need only check (1) M is irreducible, and (2) $\pi_1 M \cong G$, where G is a cocompact Kleinian group which acts freely on hyperbolic space. For in this case M and $\mathbb{H}^3/G$ are Eilenberg–MacLane spaces with isomorphic fundamental groups, and therefore homotopy equivalent.

Lemma 1.6 *The irreducibility of M is decidable.*

Proof This follows directly from classical work of Haken and the recognizability of the 3–sphere, as follows. The first step is to find a complete system of 2-spheres in M. A *complete system* is one which is guaranteed to contain enough spheres to completely decompose M into prime factors; the system may contain redundant spheres. The earliest algorithms to do this are due to Haken; more recent improvements can be found, for instance, in [10]. For each sphere S in the list, cut M along S. If the resulting manifold is connected, then S is a non-separating sphere, and so M is reducible. If $M \setminus S$ is disconnected, let M_1 and M_2 be the pieces. Each M_i has a single sphere boundary component; let

$\hat{M}_i$ be the manifold obtained by capping off this boundary component with a 3-ball. Now Rubinstein and Thompson's 3–sphere recognition algorithm may be applied to the $\hat{M}_i$, and S is a reducing sphere if and only if neither $\hat{M}_1$ or $\hat{M}_2$ is S^3. □

We can therefore assume for the remainder of the paper that M is irreducible. The question remains whether $\pi_1 M$ has a discrete faithful representation into $PSL_2\mathbb{C}$, with image which acts freely on hyperbolic space. In Section 2 we give a procedure to find a finite list of candidate representations, which must contain a discrete faithful representation if one exists. We find this list by decomposing the representation variety into irreducible components, and taking one representation from each component of the appropriate dimension. We also make some remarks on computations involving algebraic numbers, which are used heavily in the sequel.

The next two sections are devoted to determining whether or not a given candidate representation gives rise to a hyperbolic structure on M. The procedure given in Section 3 detects representations which fail to be discrete and faithful with torsion–free image. The procedure given in Section 4 detects representations which are discrete and faithful with torsion–free image, by constructing a fundamental domain for the induced action on hyperbolic space. Finally in Section 5 we prove the main theorems.

1.3 Acknowledgements

Many thanks to Andrew Casson for coming up with the idea on which this paper is based, and to my advisor Daryl Cooper for his enthusiasm and patience and for suggesting that this be written up. Thanks also to the NSF, which provided partial support while this work was being done.

2 $\overline{\mathbb{Q}}$ and some algebraic geometry in a nutshell

2.1 Computations in $\overline{\mathbb{Q}}$

The algorithms given in this paper rely heavily on the ability to do computations with algebraic numbers over the rationals. Most of what we need to use can be found in [12]. For the purposes of computation, each algebraic number is to be thought of as a minimal polynomial together with an isolating interval (or rectangle) with rational endpoints. A method for finding arbitrarily small isolating intervals for the solutions of a polynomial over $\mathbb{Q}$ can be found for instance in Section 6.2 of [3]. The method given there is for isolating real roots, but (as pointed out in an exercise) the method extends to complex roots as well. Since minimal polynomials are unique up to multiplication by an element of $\mathbb{Q}$, this representation of algebraic numbers gives us an effective test for equality.

The reader who does the exercise from [3] will note that extracting the real or imaginary part of an algebraic number is then straightforward.

For real algebraic numbers, we will need to decide the truth of statements of the form "$x < y$". Using the SIMPLE algorithm of [12], we may find a simple extension $\mathbb{Q}(\alpha)$ of the rationals which contains both x and y. The algorithm in Section 1.3 of [12] may then be used to determine the sign of $y - x$.

We will also need to be able to add, multiply and take inverses and square roots of algebraic numbers. Algorithms to perform these operations are given in [12]. The general idea is illustrated by the case of finding $\alpha\beta$, where α and β are given. First a polynomial is found which must have $\alpha\beta$ as a root. If $A(x)$ is the minimal polynomial for α and $B(x)$ is the minimal polynomial for β, then Loos shows that $\alpha\beta$ is a root of $r(x) = \mathrm{res}(y^m A(x/y), B(y))$, where res denotes the resultant. The irreducible factors of $r(x)$ are computed, and then the roots of $r(x)$ are separated from one another by isolating intervals or rectangles. Interval arithmetic is used to try to pick out the appropriate root of $r(x)$. If the isolating intervals for α and β are too big, then this can fail on the first try. In that case, the isolating intervals for α and β are shrunk until their product picks out a unique root of $r(x)$.

2.2 Finding the candidate representations

We now establish the following:

Lemma 2.1 *There is an algorithm which takes as input a triangulation of a closed 3–manifold M, and outputs a finite list of rigid representations of $\pi_1 M$ into $SL_2\overline{\mathbb{Q}}$, where $\overline{\mathbb{Q}} \subset \mathbb{C}$ is the algebraic closure of $\mathbb{Q}$. Furthermore, if M is hyperbolic, then this list contains a discrete faithful representation.*

Proof From the triangulation of M we obtain some presentation of $\pi_1 M$:

$$\pi_1 M = \langle g_1, \ldots, g_n | w_1, \ldots, w_m \rangle \tag{1}$$

We then consider the representation variety of homomorphisms

$$R = \{\rho\colon \pi_1 M \to SL_2\mathbb{C}\} \subseteq \mathbb{C}^{4n}$$

as defined for example, in [6]. Since $\pi_1 M$ is finitely presented, this variety is constructible as the zero set of a finite set of polynomials with integer coefficients. Indeed, $SL_2\mathbb{C}$ can be thought of as the variety in $\mathbb{C}^4$ defined by $\{z_1 z_4 - z_2 z_3 = 1\}$. The representation variety R is the subvariety of

$$SL_2\mathbb{C}^n = \{(z_{1,1}, \ldots, z_{4,1}, \ldots, z_{1,n}, \ldots, z_{4,n}) \mid z_{1,i}z_{4,i} - z_{2,i}z_{3,i} = 1 \ \forall i\}$$

obtained by adjoining the $4m$ polynomials coming from the matrix equations corresponding to the relations of $\pi_1 M$.

For each pair $\{g_k, g_l\}$ of generators of $\pi_1 M$, we construct the subvariety $R_{\{k,l\}} \subset R$ consisting of those representations for which the equations

$$\rho(g_k) = \begin{bmatrix} * & 1 \\ 0 & * \end{bmatrix}, \text{ and } \rho(g_l) = \begin{bmatrix} * & 0 \\ * & * \end{bmatrix}. \tag{2}$$

hold. In the notation introduced above, (2) corresponds to the equations $z_{2,k} = 1$, $z_{3,k} = 0$, and $z_{2,l} = 0$. Recall that a subgroup of $SL_2\mathbb{C}$ is said to be nonelementary if no element of $\mathbb{H}^3 \cup \partial\mathbb{H}^3$ is fixed by then entire subgroup. Notice that if ρ is any representation so that $\langle \rho(g_k), \rho(g_l)\rangle \subset SL_2\mathbb{C}$ is nonelementary, then ρ is conjugate to a representation in $R_{\{k,l\}}$.

Also observe that given some representation ρ of $\pi_1 M$ which does not send g_l to the identity, there are at most four representations conjugate to ρ which satisfy equations (2). Indeed, if ρ and $C^{-1}\rho C$ both satisfy (2) for some $C \in SL_2\mathbb{C}$, then $C\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ must be an eigenvector of $\rho(g_k)$. Likewise, $C\begin{bmatrix} 0 \\ 1 \end{bmatrix}$ must be an eigenvector of $\rho(g_l)$. Up to scaling the columns of C, there are at most four ways to arrange this. Given one of the four choices, the requirements that the determinant of C be one and that $C^{-1}\rho(g_k)C$ have a one in the upper right corner determine C up to multiplication by $\pm I$.

For each pair $\{g_k, g_l\}$, we would like to add the isolated points of $R_{\{k,l\}}$ to our list of candidate representations. Note that $R_{\{k,l\}}$ is given by a set of polynomials with integer coefficients. Let I be the ideal in $\mathbb{Q}[z_{1,1}, \ldots, z_{4,n}]$ generated by these polynomials. There exist algorithms (see, for instance [3], Theorem 8.101) to decompose this ideal into prime ideals. Proposition 9.29 of [3] gives an algorithm to determine the dimension of such an ideal. Each isolated point in $R_{\{k,l\}}$ is part of the variety determined by some zero–dimensional ideal J. Let $\{\zeta_{1,1}, \ldots \zeta_{4,n}\}$ be some such solution. Each coordinate $\zeta_{i,j}$ is a root of the monic polynomial $p(z_{i,j})$ which generates the elimination ideal $J \cap \mathbb{Q}[z_{i,j}]$. An algorithm to find this polynomial is given in Section 6.2 of [3] (see also Lemma 6.50 of [3]). Possibilities for $\zeta_{i,j}$ are determined by finding the irreducible factors of p and isolating intervals or rectangles for the various roots. There is now a finite list of possible values for each $\zeta_{i,j}$. Not every combination of choices necessarily corresponds to a representation, so we must check the original relations in $\pi_1 M$ by matrix computations over $\overline{\mathbb{Q}}$, for each of the finite number of combinations of choices. Those $4n$–tuples of choices which satisfy the relations are the candidate representations coming from $R_{\{k,l\}}$. Since there are finitely many pairs $\{g_k, g_l\}$, the process described ends up with a finite list of representations.

Suppose that M is hyperbolic. Then there is a discrete faithful representation ρ_0 of $\pi_1 M$ into $PSL_2\mathbb{C}$ which gives the hyperbolic structure on M. This representation lifts to $SL_2\mathbb{C}$ as proved in [6]. We abuse notation slightly by continuing to refer to the lifted representation as ρ_0. Note that the image $\Gamma = \rho_0(\pi_1 M) \subset SL_2\mathbb{C}$ cannot be cyclic, and that every nontrivial element is loxodromic. Therefore, there must be some pair of generators $\{g_k, g_l\}$ so that $\rho_0(g_k)$ and $\rho_0(g_l)$ are loxodromics with distinct axes. As Γ is discrete, $\langle \rho_0(g_k), \rho_0(g_l)\rangle \subset SL_2\mathbb{C}$ is nonelementary (see [2], Theorem 5.1.2). Therefore ρ_0 is conjugate to a representation ρ_0' in $R_{\{k,l\}}$. It is a well–known consequence of Mostow Rigidity that ρ_0 is unique up to conjugacy and orientation, and that nearby representations are necessarily conjugate, so ρ_0' is an isolated point in $R_{\{k,l\}}$. The discrete faithful representation ρ_0' must therefore appear in our list of candidate representations. □

Remark 2.2 If $R_{\{k,l\}}$ contains a higher dimensional curve of representations at least one of which does not send g_l to I, then it follows from [6] that M is Haken, and the question of hyperbolicity could be answered by checking whether the characteristic submanifold is empty, as pointed out in the introduction. However, the methods of the current paper still apply.

3 A procedure for rejecting representations

Suppose M is an orientable closed 3–manifold, and $\rho\colon \to SL_2\overline{\mathbb{Q}}$ is a representation as might be produced by Lemma 2.1. Suppose for the moment that ρ is discrete and faithful and that its image Γ acts freely on hyperbolic space. We immediately know a number of things about ρ and its image. For instance:

(1) ρ is irreducible (in particular, no geodesic is preserved by $\rho(\pi_1 M)$).

(2) ρ has trivial kernel.

(3) Γ contains no nontrivial elliptics or parabolics. Equivalently, no nontrivial element of Γ has trace in the interval $[-2,2]$.

(4) Given a point in hyperbolic space, the subgroup of Γ generated by elements which move that point a very small distance is abelian (in fact cyclic).

Properties (1) and (3) follow immediately from the fact that $\mathbb{H}^3/\Gamma$ is a closed 3–manifold, and that no nontrivial element of Γ has fixed points. If property (1) were to fail, then Γ would be elementary, and thus could not be cocompact. If property (3) were to fail, Γ would either not act freely (if there were elliptics) or not be cocompact (if there were parabolics). For more details see [2], especially Chapters 4 and 5. Property (2) is obvious. Property (4) follows from and is made precise by the Margulis Lemma:

Theorem 3.1 (Margulis Lemma) *There exists a small positive constant μ (eg, $\mu = \frac{.49}{16(2(4\pi/.49)^3+1)}$ works, see [1], page 107) such that the subgroup $\Gamma_\mu(M,x) \subset \pi_1(M,x)$ generated by loops based at $x \in M$ of length at most μ is abelian, for any hyperbolic three–manifold M, and any $x \in M$. In particular, if M is closed, then $\Gamma_\mu(M,x) \cong \mathbb{Z}$ or is trivial.*

Furthermore, at least one of these four properties must fail if ρ is not both discrete and faithful:

Lemma 3.2 *If $G < SL_2\mathbb{C}$ is indiscrete, consists only of loxodromics, and does not preserve any geodesic, then given $x \in \mathbb{H}^3$ and $\epsilon > 0$ there are noncommuting matrices which move x a distance less than ϵ.*

Proof For $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ a two–by–two complex matrix, define:

$$\|A\| = (|a|^2 + |b|^2 + |c|^2 + |d|^2)^{1/2}$$

The norm $\|\cdot\|$ gives rise to a metric on the space of two–by–two complex matrices defined by $\delta(A,B) = \|A-B\|$. The restriction of this metric to $SL_2\mathbb{C}$ gives the same topology as the ordinary one (see [2] for details). Therefore the hypothesis that G is indiscrete is equivalent to the existence of nontrivial elements A of G so that $\|A-I\|$ is arbitrarily small. It is shown in [2] that if $A \in SL_2\mathbb{C}$ acts on the upper half space model of hyperbolic space in the usual way, and j is the point $(0,0,1)$, then $\|A\|^2 = 2\cosh d(j, Aj)$, where d is the hyperbolic metric. A simple calculation shows that if there are nontrivial matrices $A \in G$ with $\|A-I\|$ arbitrarily small, then there must also be matrices with $\|A\|^2 - 2$ (equivalently $\cosh^{-1}(\frac{\|A\|^2}{2})$) arbitrarily small. Since the only elements of $SL_2\mathbb{C}$ which fix j (and thus have $\|A\|^2 - 2 = 0$) are elliptics whose axes pass through j, this means there must be nontrivial matrices in G with $\cosh^{-1}(\frac{\|A\|^2}{2})$ arbitrarily small but positive. We claim that we can find such matrices which do not commute.

Suppose we cannot. Then all matrices in G which move j a sufficiently small distance commute, and thus share a common axis, which we'll call γ. Let G_γ be the subgroup of G which sends γ to itself.

Rotation angle and translation length are continuous functions on $SL_2\mathbb{C}$, so we can find elements of G_γ which both translate γ an arbitrarily small distance along itself and rotate $\mathbb{H}^3$ about γ by an arbitrarily small angle.

By hypothesis, $G \setminus G_\gamma$ is nonempty. Let $C \in G \setminus G_\gamma$. The subgroup $C(G_\gamma)C^{-1}$ of G preserves the geodesic $C(\gamma) \neq \gamma$ and so its elements do not commute with those of G_γ. Since rotation angle and translation length are invariants of conjugacy class, there are elements of $C(G_\gamma)C^{-1}$ with arbitrarily small such. There are therefore elements of $C(G_\gamma)C^{-1}$ which move j as small a distance as we like. This contradicts our assumption that all matrices in G which move j a sufficiently small distance commute.

Finally let x be an arbitrary point in H^3. If g is an element of $SL_2\mathbb{C}$ which takes x to j we may apply the above argument to $g^{-1}Gg$. $\square$

The procedure for rejecting representations, outlined in Figure 1, is based on the conditions (1)–(4), so it requires that we be able to recognize when these conditions have failed.

Lemma 3.3 *There is an algorithm to decide whether or not a finite set of matrices in $SL_2\overline{\mathbb{Q}}$ share a common eigenvector.*

Proof Suppose $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in SL_2\overline{\mathbb{Q}} \setminus \{\pm I\}$.

If $c \neq 0$ then A has $\begin{bmatrix} (a \pm \sqrt{(a+d)^2 - 4})/(2c) \\ 1 \end{bmatrix}$ as eigenvectors. If $c = 0$ then A has eigenvectors $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} b/(a-d) \\ 1 \end{bmatrix}$, unless $a - d = 0$ in which case A has $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ as its sole eigenvector. Of course any multiple of an eigenvector is an eigenvector, but we have chosen specific ones precisely to remove that

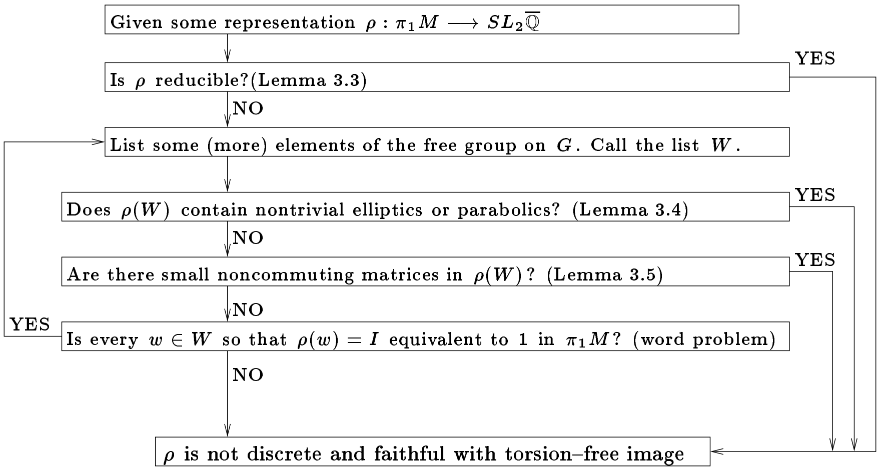

Figure 1: A procedure to reject representations. G is a finite generating set for $\pi_1 M$.

ambiguity. Since we can do arithmetic and square roots in $\overline{\mathbb{Q}}$ and compare elements of $\overline{\mathbb{Q}}$, it is possible to decide whether or not two vectors of the given form are the same. □

Lemma 3.4 *There is an algorithm to decide whether or not a finite set of matrices $S \in SL_2\overline{\mathbb{Q}}$ contains any nontrivial elliptics or parabolics.*

Proof A matrix in $SL_2\mathbb{C}$ acts elliptically or parabolically on hyperbolic space if and only if its trace is in the interval $[-2,2]$. For each $A \in S$ we can find the algebraic number trace(A). As discussed in Section 2, there are algorithms to find the imaginary part of trace(A), and to test equalities and inequalities, so it can be determined whether or not trace$(A) \in [-2,2]$. If trace$(A) = 2$, we may then check whether or not $A = I$. □

To apply the discreteness criterion effectively, we need a checkable condition on the elements of Γ. As noted in the proof of Lemma 3.2, a matrix A acting on hyperbolic space moves j a hyperbolic distance of $\cosh^{-1}(\frac{(\|A\|)^2}{2})$. If Γ is discrete, then the Margulis Lemma implies that all matrices A with $\cosh^{-1}(\frac{(\|A\|)^2}{2}) < \frac{.49}{16(2(4\pi/.49)^3+1)}$ must commute with each other. To avoid having to calculate with transcendental functions and numbers, we note that, for instance $\|A\|^2 < 2 + 2^{-58}$ implies $\cosh^{-1}(\frac{(\|A\|)^2}{2}) < \frac{.49}{16(2(4\pi/.49)^3+1)}$.

Lemma 3.5 *There is an algorithm to decide whether or not a finite set of matrices $S \in SL_2\overline{\mathbb{Q}}$ contains noncommuting elements A and B so that $\|A\| - 2 < 2^{-58}$ and $\|B\| - 2 < 2^{-58}$.*

Proof For each $A \in S$, we compute $\|A\|$ (which involves just multiplication, addition, and a square root). For each pair of matrices A, B of S so that $\|A\| - 2 < 2^{-58}$ and $\|B\| - 2 < 2^{-58}$ we check whether or not $AB - BA$ is the zero matrix. □

Now we can prove the main result of this section.

Theorem 3.6 *There exists an algorithm which takes as input a triangulated orientable closed three-manifold M, a solution to the word problem in $\pi_1 M$ and a representation $\rho\colon \pi_1 M \to SL_2\overline{\mathbb{Q}}$, and decides if the representation is* not *discrete and faithful with torsion–free image.*

Proof The idea is to detect the failure of conditions (1)–(4) listed at the beginning of this section. By Lemma 3.2, one of these conditions must fail if ρ is not discrete and faithful with torsion–free image.

We start with condition (1). Let $G = \{g_1, \ldots, g_n\}$ from the presentation (1) in Section 2.2. The representation ρ is reducible if and only if all the elements of $\rho(G)$ share a common eigenvector. This condition is checkable by Lemma 3.3.

If the representation is irreducible, we begin systematically listing elements of the free group $\mathcal{F}(G)$ in such a way that any element is eventually listed. For each finite list produced this way, we look for evidence that conditions (2)–(4) have failed for ρ. Let W be the finite list at some stage in the listing process. By performing matrix multiplications over $\overline{\mathbb{Q}}$, we can determine $\rho(w)$ for any $w \in W$. Let $P = \rho(W)$.

If condition (2) fails, then there will eventually be elements of W which are not equivalent to the identity in $\pi_1 M$, but which are in the kernel. If $\rho(w) = I$ for some $w \in W$, then we apply the solution of the word problem in $\pi_1 M$ to determine whether or not $[w]$ is the identity element in $\pi_1 M$. Thus if condition (2) fails, we will eventually discover it. Note that this is the first use of the assumption that $\pi_1 M$ has solvable word problem.

If condition (3) fails, then eventually P will contain nontrivial elliptics or parabolics. By Lemma 3.4, we can detect this.

If condition (4) fails, then there are arbitrarily small noncommuting matrices. In particular, P will eventually contain noncommuting matrices A and B with

$\|A\| - 2 < 2^{-58}$ and $\|B\| - 2 < 2^{-58}$. By Lemma 3.5, we can detect such matrices.

If ρ is discrete and faithful with torsion–free image, then none of the conditions (1)–(4) can fail, and so this procedure will continue endlessly. Otherwise, one of the four conditions will eventually be discovered to fail and the procedure will stop. $\square$

4 Recognizing a discrete faithful representation

In this section we give a procedure which takes as input a representation $\rho\colon \pi_1 M \to SL_2\overline{\mathbb{Q}}$ and decides if ρ is discrete and faithful with torsion free image. The plan of attack is to first try to build a fundamental region for the action, à la Riley in [18], and then to use the solution to the word problem to verify that ρ is injective.

It is convenient to consider both the ball and upper half–space models of hyperbolic space as subsets of the quaternions (this is a point of view taken in [2] and in [18]). Specifically we can declare $\mathbb{B}^3 = \{a + bi + cj \mid a^2 + b^2 + c^2 < 1\}$ and $\mathbb{H}^3 = \{a + bi + cj \mid c > 0\}$ and agree to always pass between them by way of the isometry $f\colon \mathbb{H}^3 \to \mathbb{B}^3$ given by $f(w) = (w - j)(w + j)^{-1}j$. A matrix $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in SL_2\mathbb{C}$ acts on $\mathbb{H}^3$ as $q \mapsto (aq + b)(cq + d)^{-1}$, and acts on $\mathbb{B}^3$ by $f \circ A \circ f^{-1}$. The action on $\mathbb{B}^3$ extends to a Möbius transformation of $\mathbb{R}^3 \cup \{\infty\}$, where $\mathbb{R}^3 = \{a + bi + cj + dk \mid d = 0\}$, and if $A \in SL_2\overline{\mathbb{Q}}$, then this action preserves the set $\overline{\mathbb{Q}}^3 \cup \{\infty\}$.

So long as A is not elliptic or $\pm I$ there is a unique Euclidean sphere in $\mathbb{R}^3$ on which the Möbius transformation acts isometrically (ie, the determinant of the Jacobian matrix is 1). We call this sphere the *isometric sphere of* A.

Definition 4.1 For a point $v \in \mathbb{R}^3$ and a sphere S with center c and radius r we say v is *inside* S if $\|v - c\| < r$, v is *outside* S if $\|v - c\| > r$, and v is *on* S if $\|v - c\| = r$. Note that if r and the coordinates of the points c and v are all algebraic, then each of the three conditions is decidable.

Note that since the intersection of an isometric sphere with $\mathbb{B}^3$ is a hyperbolic plane in $\mathbb{B}^3$, the intersection of the interior or exterior of an isometric sphere with $\mathbb{B}^3$ is convex with respect to the hyperbolic metric.

The isometric spheres are not intrinsic to the hyperbolic metric in the sense that, say, the axis of a loxodromic is intrinsic. However, they give us a convenient way to construct a particular fundamental domain for a discrete action. In the case that no group element acts as a nontrivial elliptic, the closure of the set of points in $\mathbb{B}^3$ which are outside every isometric sphere is a fundamental domain, sometimes called the *Ford domain* (though this term is sometimes used differently). In fact, this set is precisely the Dirichlet domain centered at 0 (see [14], page 45). This is the fundamental domain which we will construct for a discrete action.

According to Proposition 1.3 of [14], the isometric sphere of an isometry $\phi\colon \mathbb{B}^3 \to \mathbb{B}^3$ has Euclidean center $\frac{\phi^{-1}(0)}{|\phi^{-1}(0)|^2}$ and Euclidean radius $(\frac{1}{\phi^{-1}(0)^2} - 1)^{1/2}$. In particular, we have:

Lemma 4.2 *If S_A is the isometric sphere of $A \in SL_2\overline{\mathbb{Q}}$ acting on $\mathbb{B}^3$, and S_A has radius r_A and center $c_A = (x_A, y_A, z_A)$, then r_A, x_A, y_A and z_A are algebraic numbers which can be computed from the entries of the matrix A.*

Let W be some finite list of words in $\mathcal{F}(G)$, closed under inverses, and let $P = \rho(W) \setminus \{\pm I\}$. We say that P *satisfies condition NE* if no nontrivial element of P is an elliptic, and no two nontrivial elements of P have the same isometric sphere.

Lemma 4.3 *There is an algorithm to decide whether or not P satisfies condition NE.*

Proof This follows from our ability to compute traces and from Lemma 4.2. $\square$

If P satisfies condition NE, let D be the closure in $\mathbb{B}^3$ of the intersection of the exteriors of the isometric spheres S_A for $A \in P$. Then D is convex with respect to the hyperbolic metric, and bounded by finitely many (possibly noncompact or infinite area) hyperbolic polygons, each of which is a subset of some isometric sphere. Define a *vertex* of D to be a point in D which is contained in at least three distinct isometric spheres. Let K be the intersection of the convex hull of the vertices of D with ∂D. Figure 2 gives an idea what K and D might look like one dimension down (in dimension 2 vertices need only be contained in two distinct isometric spheres). K can be thought of as the "compact part" of the boundary of D. In particular, if D is compact, then $K = \partial D$. K has an obvious cell decomposition. In particular, the vertices of K are precisely

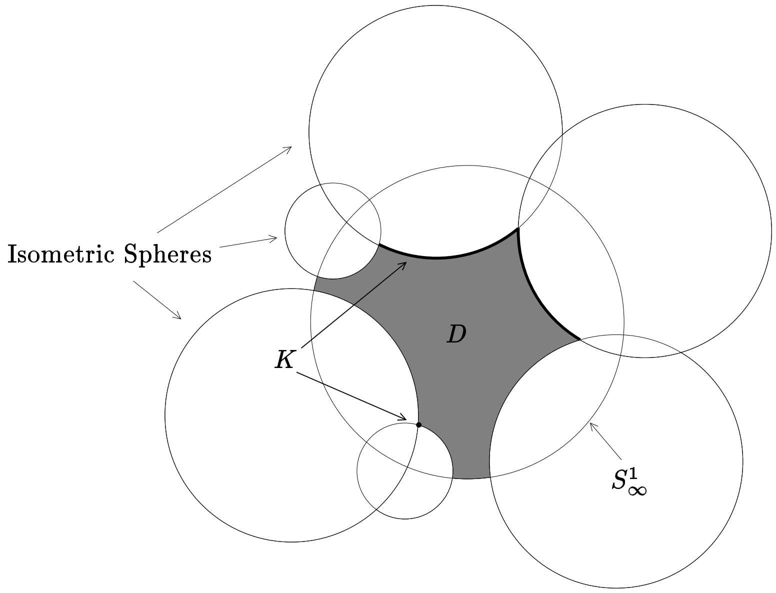

Figure 2: This is what K and D might look like in dimension 2. In this example K has four vertices and two edges.

the vertices of D. If a pair of vertices are contained in two or more distinct isometric spheres, then their convex hull forms an edge of K. The 2–cells of K are the convex hulls of maximal collections of at least three vertices, subject to the condition that the convex hull is contained in K. Every cell in K is the convex hull of its vertices, and is thus completely determined by them.

Lemma 4.4 *If P satisfies condition NE, there is an algorithm to construct K. By constructing K we mean computing the locations of the vertices of K, and specifying the edges and faces of K as subsets of the vertex set of K.*

Proof Given three spheres in Euclidean space, it is not hard to write down a set of inequalities in terms of their centers and radii which are satisfied if and only if the spheres intersect in a pair of distinct points. Moreover, these inequalities involve only arithmetic and square roots, and so may be checked algorithmically if the centers and radii are given by algebraic numbers. If the spheres intersect in a pair of points, then the coordinates of these points can be explicitly written down (again using only arithmetic and square roots) in terms of the centers and radii of the spheres.

Therefore, for each triple of isometric spheres of elements of P it is possible to decide if they intersect in $\mathbb{B}^3 \subset \mathbb{R}^3$, and if so to write down a triple of algebraic numbers which are the coordinates of their intersection. These triples are potential vertices of K.

Let $\mathcal{V}$ be the set of potential vertices of K. Since every vertex of K is contained in at least three distinct isometric spheres, every vertex of K is in $\mathcal{V}$. However $\mathcal{V}$ may contain vertices not in D. For each $A \in P$, $v \in \mathcal{V}$, we check that v is outside or on the sphere S_A. If v is inside S_A, v is outside D, and so it is discarded. After discarding these vertices, we continue to refer to the smaller set of vertices as $\mathcal{V}$. After all sphere–vertex pairs have been checked, this set is equal to the zero–skeleton of K.

Now we find the edges of K by looking for pairs of spheres whose intersection contains two vertices. For each pair of vertices $\{v, w\}$ in $\mathcal{V}$, we can construct the list of matrices A in P so that v and w are both on S_A. There is an edge connecting v to w in K if and only if this list contains at least two elements. Let $\mathcal{E}$ be the set of unordered pairs of vertices contained in at least two isometric spheres. Clearly every edge of K is represented by some element of $\mathcal{E}$. Conversely, suppose $\{v, w\} \subset S_A \cap S_B$ for some $A, B \in P$. Since D is convex, the geodesic segment between v and w is contained in D. Since $S_A \cap S_B$ is convex, this segment is in fact contained in the boundary of D, and is therefore part of K. Thus we can construct the 1–skeleton of K.

It remains only to find the 2–cells of K. For each $A \in P$, the set $\mathcal{V} \cap S_A$ can be constructed. There is a 2–cell of K formed from a part of S_A if and only if the set $\mathcal{V} \cap S_A$ contains at least three vertices and these vertices are linked up by edges in $\mathcal{E}$ to form a circuit. These conditions can be checked on S_A for each $A \in P$, so we can produce a list of faces $\mathcal{F}$. This completes the combinatorial data needed to construct K. □

Remark 4.5 (Notation) It is no extra work in the above to keep track of which matrix corresponds to which face. Since each isometric sphere contributes at most one face to $\mathcal{F}$ and since P satisfies condition NE, we may unambiguously refer to a face as F_A if it is contained in the isometric sphere of A.

Lemma 4.6 *There is an algorithm to decide whether or not K is a 2–sphere.*

Proof To check that K is connected, it suffices to check that any pair of vertices in $\mathcal{V}$ can be connected by a sequence of edges. The Euler characteristic of K is equal to $|\mathcal{F}| - |\mathcal{E}| + |\mathcal{V}|$. □

Note that if K is a 2–sphere, then D is a compact finite sided hyperbolic polyhedron with $K = \partial D$. If D is a fundamental domain for a discrete action of some group on hyperbolic space, then we should have $A(F_A) = F_{A^{-1}}$ for every $F_A \in \mathcal{F}$. If this is the case, then K can be given the additional structure of a polyhedron with face identifications.

Lemma 4.7 *If the complex K produced by Lemma 4.4 is a 2–sphere, then there is an algorithm to check whether or not $A(F_A) = F_{A^{-1}}$ for every $F_A \in \mathcal{F}$.*

Proof We first check, for each $F_A \in \mathcal{F}$, that we also have $F_{A^{-1}} \in \mathcal{F}$. Since faces are uniquely determined by their vertices, it then suffices to check that A sends the vertices of F_A to those of $F_{A^{-1}}$. Since we have kept track of the coordinates of the vertices, and can explicitly compute the action of A on points in $\mathbb{B}^3$ with algebraic coordinates, this is straightforward. □

The question of when a polyhedron with face identifications is the fundamental domain of a discrete action consistent with those face identifications is answered by the Poincaré Polyhedron Theorem. We use the theorem in the following specialized form (more general versions can be found in the standard references [13] and [21]):

Theorem 4.8 (Poincaré Polyhedron Theorem) *Let D be a compact polyhedron in hyperbolic space, together with face pairings which are orientation-preserving isometries generating the group Γ. Suppose the angle sum about every image of an edge in the identified polyhedron is 2π. Then:*

(1) *D is a fundamental polyhedron for the (fixed–point free) action of Γ, and*

(2) *Γ has a presentation as an abstract group with the generators equal to the face pairings and the relations precisely the edge cycles.*

Remark 4.9 Theorem 4.8 gives a sufficient condition for a compact convex polyhedron with face pairings to be the fundamental domain for a discrete cocompact action. If we also require that the action be fixed–point free (so that the quotient is a hyperbolic 3–manifold), then the sufficient condition is also clearly necessary.

Lemma 4.10 *If the complex K produced by Lemma 4.4 is a 2–sphere which satisfies $A(F_A) = F_{A^{-1}}$ for every $F_A \in \mathcal{F}$, then there is an algorithm to decide whether or not K is the boundary of a fundamental domain for a discrete fixed–point free action on hyperbolic space, generated by $\{A \in P \mid F_A \in \mathcal{F}\}$.*

Proof We use a combination of exact and numerical computations to check whether the conditions of Theorem 4.8 are satisfied at each edge. Specifically, let e_1 be some edge of K. Adjacent to e_1 are precisely two faces, F_A and F_B. Let $F_{A_1} = F_A$ be one of the two faces. Let $e_2 = A_1(e_1)$. By hypothesis, e_2 is an edge of K. The edge e_2 is adjacent to two faces, one of which is $A_1(F_{A_1})$. The other face is F_{A_2} for some $A_2 \in P$. We inductively define e_{i+1} to be $A_i(F_{A_i})$ and $F_{A_{i+1}}$ to be the face adjacent to $A_i(F_{A_i})$ across e_{i+1}. Let k be the first integer for which $e_k = e_1$, the edge we started with, in which case $A_{k-1}(F_{A_{k-1}}) = F_B$ and $F_{A_k} = F_B$.

If $\{e_1, \ldots, e_{k-1}\}$ satisfies the edge condition, then, in particular, $A_{k-1} \circ \cdots \circ A_1 = I$. If the exact calculation shows that this is the case, then the angle sum is some integer multiple of 2π.

Suppose S_1 and S_2 are two spheres whose intersection contains an edge. Then their dihedral angle is given by $\cos^{-1}(\frac{\|c_1 - c_2\|^2 - {r_1}^2 - {r_2}^2}{2 r_1 r_2}) \in (0, \pi)$, if S_i is the sphere of radius r_i centered at c_i. In the case where r_i and the coordinates of c_i are given by algebraic numbers (or any numbers which can be determined to any desired accuracy) it is not hard to show that the angle can then be determined numerically to any desired accuracy. We can therefore do a calculation of the angle sum around an edge with tolerance less than, say $\pi/2$. Together with the exact calculation, this tells us whether the edge condition is satisfied. As there are a finite number of edges in the complex K, it will take a finite amount of time to check them all. □

Remark 4.11 The proof of the preceding lemma can be modified to deal with the case of an arbitrary polyhedron with identifications, provided that all the identifying isometries are given as matrices in $SL_2\overline{\mathbb{Q}}$, and all vertices lie inside S^2_∞ and have algebraic coordinates. The version given is sufficient for our purposes, however.

If K as constructed in Lemma 4.4 is the boundary of a fundamental domain, we will refer to the group generated by $\{A \in P \mid F_A \subset K\}$ as Γ_K. Γ_K is clearly a subgroup of $\rho(\pi_1 M)$, and the quotient of hyperbolic space by the action of Γ_K is a closed hyperbolic 3–manifold.

Remark 4.12 It is possible that Γ_K is a proper subgroup of $\rho(\pi_1 M)$. In the case that $\rho(\pi_1 M)$ is discrete, then the quotient of hyperbolic space by Γ_K is a finite cover of the quotient of hyperbolic space by $\rho(\pi_1 M)$.

Lemma 4.13 *If K, constructed as above from a representation $\rho\colon \pi_1 M \to SL_2\overline{\mathbb{Q}}$, is the boundary of a fundamental domain D of a discrete fixed–point* *free action on hyperbolic space, then there is an algorithm to decide whether or not $\Gamma_K = \rho(\pi_1 M)$, where Γ_K is the group generated by the elements of $\rho(\pi_1 M)$ which pair the faces of K.*

Proof For each generator g of $\pi_1 M$, we wish to check whether $\rho(g)$ is in Γ_K. Let $\{T_1, \ldots, T_s\}$ be the generators of Γ_K, which are also the face pairings of D. Each face of K is part of the isometric sphere S_T for some $T \in \{T_1, \ldots, T_s\} \cup \{T_1^{-1}, \ldots, T_s^{-1}\}$.

Note that $\mathbf{0} = (0,0,0)$ is in the interior of D. Since hyperbolic space is tiled by copies of D, $\rho(g)(\mathbf{0})$ is in $\gamma(D)$ for some $\gamma \in \Gamma_K$. But the following are clearly equivalent:

(1) $\rho(g)(\mathbf{0}) \in \gamma(D)$

(2) $\gamma^{-1}\rho(g)(\mathbf{0}) \in D$

(3) $\gamma^{-1}\rho(g)(\mathbf{0})$ is outside or on S_T for all $T \in \{T_1, \ldots, T_s\} \cup \{T_1^{-1}, \ldots, T_s^{-1}\}$.

From our earlier remarks, this last statement is clearly checkable. If $\rho(g) \in \Gamma_K$ then $\rho(g) = \gamma$ whenever the statement is true.

We can systematically list all possible words in the free group $\mathcal{F}(T_1, \ldots, T_s)$ and for each word w in the list check to see whether condition (3) above holds for $w^{-1}\rho(g)(\mathbf{0})$. We must eventually find such a w, since the copies of D tile hyperbolic space. When this w is found, test whether $w^{-1}\rho(g) = I$. If so, then $\rho(g) \in \Gamma_K$, and we go on to the next generator. If every generator has image in Γ_K, then $\Gamma_K = \rho(\pi_1 M)$. Otherwise, Γ_K is a proper subgroup of $\rho(\pi_1 M)$. □

Lemma 4.14 *Suppose $\rho\colon \pi_1 M \to SL_2\overline{\mathbb{Q}}$, and K is a complex constructed as above from a finite subset of $\rho(\pi_1 M)$, and that $\Gamma_K = \pi_1 M$. If there is an algorithm to solve the word problem in $\pi_1 M$, then there is an algorithm to decide whether or not ρ is faithful.*

Proof We can now think of ρ as an epimorphism of abstract groups

$$\rho\colon \langle g_1, \ldots, g_n | r_1, \ldots, r_m \rangle \longrightarrow \langle T_1, \ldots, T_s | c_1, \ldots, c_r \rangle$$

where the c_i are words coming from the edge cycles. Recall that we started with some finite set W of words in the free group $\langle g_1, \ldots, \ldots, g_n \rangle$. The T_i correspond to the matrices in $P = \rho(W)$ which pair the faces of the polyhedron K, and generate Γ_K. To check whether or not ρ is faithful, we attempt to construct an inverse map. For each T_i pick some w_i in the finite set $\rho^{-1}(T_i) \cap W$. The w_i thus chosen determine a homomorphism from the free group $\langle T_1, \ldots, T_s \rangle$ to

$\pi_1 M$. We can check whether this map factors through Γ_K by using the word problem algorithm to check that each c_i is sent to the identity element of $\pi_1 M$. If the map does not factor, then ρ does not have an inverse and therefore must have nontrivial kernel.

Otherwise, we have constructed a homomorphism

$$\phi\colon \langle T_1, \dots, T_s | c_1, \dots, c_r \rangle \longrightarrow \pi_1 M$$

ϕ from $\langle T_1, \dots, T_s | c_1, \dots, c_r \rangle$ to $\pi_1 M$ and we have $\rho \circ \phi = \mathrm{Id}_{\Gamma_K}$ by construction. To determine whether $\phi \circ \rho = \mathrm{Id}_{\pi_1 M}$ the procedure checks that for each generator $\phi \circ \rho(g_i) = g_i$, again using the word problem algorithm. If any $\phi \circ \rho(g_i) \neq g_i$, then $\rho((\phi \circ \rho)(g_i) g_i^{-1}) = \rho(g_i)\rho(g_i^{-1}) = I$ implies that ρ has kernel. Otherwise, ρ is an isomorphism onto Γ_K. □

Theorem 4.15 *There exists an algorithm which takes as input a triangulated orientable closed three-manifold M, a solution to the word problem in $\pi_1 M$ and a representation $\rho\colon \pi_1 M \to SL_2\overline{\mathbb{Q}}$, and decides if the representation is discrete and faithful with torsion–free image.*

Proof Let $G = \{g_1, \dots, g_n\}$ from the presentation (1) in Section 2.2. We begin systematically listing elements of the free group $\mathcal{F}(G)$ in such a way that an element and its inverse are always added at the same time. At each stage of the listing, we have some finite list W of words and a list $P = \rho(W)\{\pm I\}$ of matrices in $SL_2\overline{\mathbb{Q}}$. From this data we attempt to build a fundamental domain.

Taken together, Lemmas 4.2, 4.3, 4.4, 4.6, 4.7, 4.10 and 4.13 give an algorithm which either constructs or fails to construct a fundamental domain for the action of Γ, given some finite $P \subset \Gamma$ which is closed under inverses. This algorithm is outlined in Figure 3.

If this algorithm finds a fundamental domain for the action of Γ, then it follows that Γ acts as a discrete fixed point free group of isometries of hyperbolic space. According to Lemma 4.14, there is then an algorithm to determine whether ρ is faithful.

Let $\rho\colon \pi_1 M \to SL_2\overline{\mathbb{Q}}$ be a discrete faithful representation with torsion–free image Γ. We must show that the algorithm we have described eventually discovers a fundamental domain for the action of Γ. Note first that any subset of Γ satisfies condition NE. For suppose that A and B are hyperbolics in Γ which have the same isometric sphere $S_A = S_B$. Then AB^{-1} fixes $\mathbf{0}$, and so must be trivial or elliptic. But Γ contains no elliptics, so we must have $A = B$.

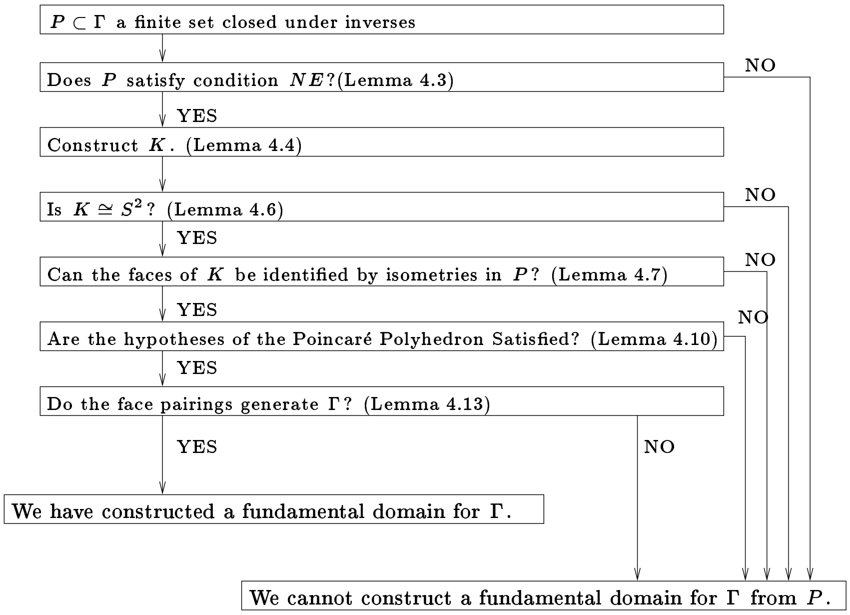

Figure 3: An algorithm which tries to construct a fundamental domain for $\Gamma < SL_2\overline{\mathbb{Q}}$ from a finite subset of Γ. K is the "compact part" of the boundary of D, if D is the region outside all the isometric spheres of elements of P.

The intersection of the exteriors of *all* the isometric spheres of elements of Γ is the Dirichlet domain for Γ centered at $\mathbf{0}$. Theorem 1.5 implies that the quotient of hyperbolic space by Γ is homeomorphic to M, so it is a closed 3–manifold. In particular, Γ is geometrically finite, so the Dirichlet domain centered at $\mathbf{0}$ has finitely many sides. Each of these sides is part of the isometric sphere of a unique loxodromic element of Γ. There is therefore a finite subset of Γ containing all the matrices whose isometric spheres form faces of the Dirichlet domain. Eventually our list P will contain all these matrices, and we suppose for the remainder of the proof that it does. But now the complex K constructed in Lemma 4.4 is precisely the boundary of the Dirichlet domain centered at $\mathbf{0}$, which is homeomorphic to a sphere. Face pairings are given by the isometries whose isometric spheres form the faces, and the algorithm of Lemma 4.7 will verify this. The Dirichlet domain satisfies the hypotheses of the Poincaré Polyhedron Theorem, as will be verified by the algorithm of 4.10. The face identifications of the Dirichlet domain for an action generate that action, so the algorithm of 4.13 will correctly determine that $\Gamma_K = \pi_1 M$. Thus if P is large enough, the algorithm of Figure 3 will eventually produce a fundamental domain for the action of Γ. □

5 Main theorems

The main theorem of the paper now follows easily from Theorem 3.6 and 4.15.

Theorem 5.1 *There exists an algorithm which will, given a triangulated orientable closed three-manifold M and a solution to the word problem in $\pi_1 M$, decide whether or not M has a hyperbolic structure.*

Proof By Theorem 1.5 M admits a hyperbolic structure if and only if

(1) M is irreducible, and
(2) there is an injective homomorphism $\rho\colon \pi_1 M \to SL_2\mathbb{C}$ so that $\rho(\pi_1 M)$ acts freely and cocompactly on $\mathbb{H}^3$.

By Lemma 1.6 we may suppose that M is irreducible. Lemma 2.1 gives a way to construct a finite list $\mathcal{L}$ of representations of $\pi_1 M$ into $SL_2\overline{\mathbb{Q}} \subset SL_2\mathbb{C}$, with the property that if a representation satisfying condition 2 exists, then it is in $\mathcal{L}$.

Theorems 4.15 and 3.6 together give a way to decide whether a particular representation from $\mathcal{L}$ satisfies condition 2 or not. The manifold M admits a hyperbolic structure if and only if some representation in $\mathcal{L}$ satisfies the condition. □

Theorem 5.2 *If M is a triangulated closed orientable 3–manifold which has a hyperbolic structure, then the hyperbolic structure is algorithmically constructible.*

Proof Note that by Remark 1.2, it is sufficient to find a triangulated hyperbolic manifold which is homeomorphic to M. If it is known that M is hyperbolic, it follows that the fundamental group is automatic. There are algorithms to find the automatic structure once one is known to exist [7]. This automatic structure immediately provides a procedure for solving the word problem. We can then use the algorithms of Sections 2 and 4 to find a particular representation of $\pi_1 M$ which is discrete and faithful, and a fundamental domain for the action of $\pi_1 M$ on hyperbolic space through that representation. The fundamental polyhedron with face–identifications give a cell–structure for $N = \mathbb{H}^3/\pi_1 M$ which is easily turned into a triangulation. □

References

[1] **Werner Ballman**, **Mikhael Gromov**, **Victor Schroeder**, *Manifolds of Non-positive Curvature*, Birkhäuser, Boston (1985)

[2] **Alan F Beardon**, *The Geometry of Discrete Groups*, Springer–Verlag, New York (1983)

[3] **Thomas Becker**, **Volker Weisspfenning**, *Gröbner Bases: A Computational Approach to Commutative Algebra*, Springer–Verlag, New York (1993)

[4] **R H Bing**, *Locally tame sets are tame*, Annals of Mathematics 69 (1954) 145–158

[5] **Daryl Cooper**, **Craig Hodgson**, **Steven Kerckhoff**, *Thurston's Orbifold Theorem*, to appear

[6] **Marc Culler**, **Peter B Shalen**, *Varieties of group representations and splittings of 3–manifolds*, Annals of Mathematics 117 (1983) 109–146

[7] **David B A Epstein**, **James W Cannon**, **Derek F Holt**, **Silvio V F Levy**, **Michael S Paterson**, **William P Thurston**, *Word Processing in Groups*, Jones and Bartlett Publishers, Boston (1992)

[8] **David Gabai**, **G Robert Meyerhoff**, **Nathaniel Thurston**, *Homotopy Hyperbolic 3-Manifolds are Hyperbolic*, MSRI 1996-058

[9] **Geoffrey Hemion**, *On the classification of homeomorphisms of* 2*–manifolds and the classification of* 3*–manifolds*, Acta Mathematica 142 (1979) 123–155

[10] **William Jaco**, **Jeffrey L Tollefson**, *Algorithms for the complete decomposition of a closed 3-manifold*, Illinois Journal of Mathematics 39 (1995) 358–406

[11] **Mark Lackenby**, *Word hyperbolic Dehn surgery*, Inventiones Mathematicae 140 (2000) 243–282

[12] **Rüdiger Loos**, *Computing in Algebraic Extensions*, from: "Computer Algebra: Symbolic and Algebraic Computation", (Bruno Buchberger, George Edwin Collins, Rüdiger Loos, Rudolf Albrecht, editors), Springer–Verlag, New York (1983) 173–188

[13] **Bernard Maskit**, *On Poincaré's Theorem for Fundamental Polygons*, Advances in Mathematics 7 (1971) 219–230

[14] **Katsuhiko Matsuzaki**, **Masahiko Taniguchi**, *Hyperbolic Manifolds and Kleinian Groups*, Oxford University Press, Oxford (1998)

[15] **Edwin E Moise**, *Affine structures in* 3*–manifolds, V: the triangulation theorem and the Hauptvermutung*, Annals of Mathematics 72 (1952) 96–114

[16] **G A Niblo**, **L D Reeves**, *The geometry of cube complexes and the complexity of their fundamental groups*, Topology 37 (1998) 621–633

[17] **Carlo Petronio**, **Jeffrey R Weeks**, *Partially flat ideal triangulations of cusped* 3*–manifolds*, Osaka Journal of Mathematics 37 (2000) 453–466

[18] **Robert Riley**, *Applications of a computer implementation of Poincaré's theorem on fundamental polyhedra*, Mathematics of Computation 40 (1983) 607–632

[19] **Joachim Hyam Rubinstein**, *An algorithm to recognize the* 3*–sphere*, from: "Proceedings of the International Congress of Mathematicians", Birkhäuser, Basel (1995) 601–611

[20] **Peter Scott**, *The geometries of 3-manifolds*, Bulletin of the London Mathematical Society 15 (1983) 401–487

[21] **Herbert Seifert**, *Komplexe mit Seitenzuordnung*, Göttingen Nachrichten (1975) 49–80

[22] **Z Sela**, *Topics in 3-manifolds*, Ph.D. thesis, Hebrew University (1991)

[23] **Z Sela**, *The isomorphism problem for hyperbolic groups I*, Annals of Mathematics 141 (1995) 217–283

[24] **Andrew Skinner**, *The word problem in a class of non–Haken 3-manifolds*, Topology 33 (1994) 215–239

[25] **Michelle Stocking**, *Almost normal surfaces in 3-manifolds*, Ph.D. thesis, U. C. Davis (1996)

[26] **Abigail Thompson**, *Thin position and the recognition problem for* S^3, Mathematical Research Letters 1 (1994) 613–630

[27] **William P Thurston**, *Three–dimensional manifolds Kleinian groups and hyperbolic geometry*, Bulletin of the American Mathematical Society 6 (1982) 357–381

[28] **Friedhelm Waldhausen**, *The word problem in fundamental groups of sufficiently large irreducible 3–manifolds*, Annals of Mathematics (2) 88 (1968) 272–280

[29] **Friedhelm Waldhausen**, *Recent results on sufficiently large* 3*–manifolds*, from: "Algebraic and geometric topology (Proceedings of the Symposium in Pure Mathematics of the American Mathematical Society)", Proceedings of Symposia in Pure Mathematics XXXII 2, American Mathematical Society, Providence, R. I. (1978) 21–38