The homeomorphism problem for closed 3–manifolds

PETER SCOTT

HAMISH SHORT

We give a geometric approach to an algorithm for deciding whether two hyperbolic 3–manifolds are homeomorphic. We also give an algebraic approach to the homeomorphism problem for geometric, but nonhyperbolic, 3–manifolds.

1 Introduction

The homeomorphism problem for closed orientable triangulated 3–manifolds has been studied for many years, with partial results by many authors. The work of Perelman proving Thurston's geometrization conjecture [17; 18; 19] finally allowed a complete solution for irreducible such manifolds, which has been described by Jaco in his Beijing Lectures [10] and in the first chapter of the book by Bessières, Besson, Maillot, Boileau and Porti [3]. The work of many previous authors is put together, and different algorithms are used to deal with the Haken case, to find the JSJ decomposition and to deal with different geometries. In the case of two closed hyperbolic 3–manifolds M_1 and M_2, Mostow rigidity tells us that M_1 and M_2 are homeomorphic if and only if their fundamental groups are isomorphic. Thus the homeomorphism problem for M_1 and M_2 can be solved by appealing to Sela's solution [22] of the isomorphism problem for torsion–free word hyperbolic groups. The initial aim of this paper was to give a more geometric approach to the homeomorphism problem in this case, which avoids quoting Sela's work. But in addition, we also give a more algebraic approach to some other parts of the homeomorphism problem, though the geometric results of Jaco and Oertel [11] are still needed for the existence of incompressible surfaces. We are not claiming that the algorithms we present are in any way superior to those referred to by Jaco and by Bessières et al. Our aim is simply to increase the range of applicable algorithms from which to choose. In this paper, we will mostly consider closed, orientable and irreducible 3–manifolds. It is known that given two triangulated closed orientable 3–manifolds M_1 and M_2, there is an algorithm to find the geometric structures on the geometric pieces of each, and then there is an algorithm to decide whether or not the pieces of M_1 are homeomorphic to the pieces of M_2 (see eg [14, 1.4]). This algorithm

is described in [3, 1.4.1], where it is claimed that this solves the homeomorphism problem for triangulated closed orientable 3–manifolds. (A similar claim is made by Matveev in his book [15, page 214], using his method of spines for Haken manifolds.) It was pointed out to us by Henry Wilton, that there remains an orientation problem when considering connected sums. If M and N denote two closed orientable 3–manifolds and $\overline{N}$ denotes N with the opposite orientation, it is not clear how to decide whether the connected sum $M \# \overline{N}$ is homeomorphic to $M \# N$ (though the geometric pieces are clearly homeomorphic).

In [14], Manning gave an algorithm to decide if there exists a hyperbolic structure on a closed orientable 3–manifold given by a triangulation. If there is such a structure, then Manning's algorithm constructs a finite sided polyhedral fundamental region in hyperbolic 3–space. In Section 2 of this paper we use Manning's work to give a new algorithm to decide whether or not two closed 3–manifolds M_1 and M_2 are homeomorphic, when they are given by triangulations and known to be hyperbolic. Manning's algorithm will construct a finite sided polyhedral fundamental region P_i in $\mathbb{H}^3$ for each manifold M_i, and we show how to estimate how many copies of P_2 must be glued together to contain P_1 and vice versa. This allows one to bound the number of maps from the generators of $\pi_1(M_1)$ to $\pi_1(M_2)$ which might give rise to an isomorphism, and vice versa, enabling one to check each such map in turn and decide whether or not it is an isomorphism.

In Section 3 of this paper, we take a more algebraic point of view. The main ingredients are the existence of a biautomatic structure on the fundamental groups of most geometric 3–manifolds and the fact that Perelman's work essentially reduces the homeomorphism problem for closed orientable irreducible 3–manifolds to the isomorphism problem for their fundamental groups (with the notable exception of lens spaces).

Given a triangulated closed orientable 3–manifold, Jaco and Tollefson [12, Section 7] give an algorithm producing a decomposition into irreducible pieces. When the manifold is Haken, they give [12, Section 8] algorithms for finding its JSJ decomposition and finding the Seifert invariants for each Seifert piece. Thus for a Haken closed Seifert fibre space, their algorithm will yield all the Seifert invariants and hence a geometric structure. Now suppose M is a closed orientable 3–manifold which is known to satisfy Thurston's geometrization conjecture. In [14, 1.4], Manning combined the algorithms in Jaco and Tollefson [12] with his algorithm for finding hyperbolic structures to obtain an algorithm which finds the geometric structures on all the pieces of M. In this section we describe a somewhat simpler algorithm, with substantial algebraic ingredients, which can answer a somewhat simpler question. Namely, if M is a closed orientable 3–manifold which is known to be geometric, then our algorithm can decide on which geometry M is modelled. To answer this simpler question, we do not need

Manning's algorithm for hyperbolic manifolds, and our procedure differs from that of Manning in the other cases. One natural application of our algorithm would be when M is orientable and irreducible and not Haken, as then Perelman's work implies that M must be geometric. Our approach also works in the Haken case and is rather different from that in [12]. Our algorithm starts by finding a biautomatic structure for the fundamental group of M, when such a structure exists. As for many algorithms for 3–manifolds, the main topological ingredient is an algorithm due to Jaco and Oertel [11] to decide whether a triangulated 3–manifold is Haken, and if so, to produce embedded incompressible surfaces.

2 A new algorithm to decide if two hyperbolic 3–manifolds are homeomorphic

Let M_1 and M_2 be closed 3–manifolds each given by a triangulation and known to be hyperbolic. It is straightforward to use these triangulations to write down finite presentations for $\pi_1(M_1)$ and $\pi_1(M_2)$. In order to apply Manning's algorithm in [14], we first need an algorithm for solving the word problem in each $\pi_1(M_i)$. Since M_i is known to be hyperbolic, its fundamental group is word hyperbolic and so biautomatic. Then, the algorithm of [6, 3.4.1] (extended to the biautomatic case, as stated in Gersten and the second author [8, after Lemma 8.2]) produces a biautomatic structure. Alternatively Papasoglu [16] gives an algorithm for calculating the hyperbolicity constant δ of the presentation, and with that a solution to the word problem is easily built as in Alonso, Brady, Cooper, Ferlini, Lustig, Mihalik, Shapiro and the second author [1]. It follows that there is an algorithm which can be applied to the given presentation for $\pi_1(M_i)$ which will find the biautomatic structure. (We note that Chapter 5 of Epstein, Cannon, Holt, Levy, Paterson and Thurston [6] gives the algorithm to find automatic structures; the additional languages and axioms needed to extend to biautomatic structures are easily added.) Once this is done, there is an algorithm for solving the word problem in $\pi_1(M_i)$ (alternatively one can use the fact that the groups are residually finite). Now, for each M_i, Manning's algorithm in [14] constructs a convex finite sided polyhedral fundamental region P_i in $\mathbb{H}^3$. His algorithm describes the vertices, edges and faces of P_i. In addition, the algorithm describes the face pairing isometries needed to recover M_i from P_i. It is now straightforward to write down a new presentation for $\pi_1(M_i)$ with the face pairing isometries of P_i as generators. In what follows we will use these presentations rather than the presentations obtained from the initial triangulations.

Now we consider the tiling of $\mathbb{H}^3$ by translates of P_2. Given a union X of such translates, we let star(X) denote the union of all translates of P_2 which meet X in

at least one point, and let $\text{star}^n(X)$ denote the result of applying the star operation n times to X. Thus $\text{star}^2(X) = \text{star}(\text{star}(X))$. It may sometimes be convenient to write $\text{star}^0(X) = X$. We consider the sequence $\text{star}^n(P_2)$, $n \geq 0$, of subsets of $\mathbb{H}^3$. As the union of these subsets equals $\mathbb{H}^3$, their diameters must tend to infinity. We will refine this obvious fact in the following way. We will show in Lemma 2.1 below that there is an algorithm to find a positive number R, such that $\text{star}^n(P_2)$ contains the metric ball $B(P_2, nR)$.

Assuming this lemma for the moment, we now proceed as follows. If M_1 and M_2 are homeomorphic, Mostow's rigidity theorem implies that they must be isometric with their hyperbolic metrics. This yields an isometry φ from $\mathbb{H}^3$ tiled by translates of P_1 to $\mathbb{H}^3$ tiled by translates of P_2. By composing with the action of an element of $\pi_1(M_2)$ on $\mathbb{H}^3$, we can suppose that $\varphi(P_1)$ meets P_2. Let d_1 denote the diameter of P_1. Then $\varphi(P_1)$ must be contained in the metric ball $B(P_2, d_1)$. If n is an integer such that $nR > d_1$, it follows that $\varphi(P_1)$ is contained in $\text{star}^n(P_2)$. Let α be a face pairing isometry of P_1. Then the isometry $\varphi\alpha\varphi^{-1}$ of $\mathbb{H}^3$ pairs faces of $\varphi(P_1)$. As α lies in $\pi_1(M_1)$, the isometry $\varphi\alpha\varphi^{-1}$ lies in $\pi_1(M_2)$, and so preserves the tiling by translates of P_2. Hence $\varphi\alpha\varphi^{-1}$ must send a certain translate of P_2 which is contained in $\text{star}^n(P_2)$ to another such. Pick a path in $\text{star}^n(P_2)$ which joins the interiors of these two translates of P_2, does not meet any edges and is transverse to the faces, and let N denote the number of times this path meets a face. Then $\varphi\alpha\varphi^{-1}$ can be written as a word in the face pairing generators of $\pi_1(M_2)$ of length N.

Let k_1 denote the maximum number of translates of P_2 around an edge, and k_2 denote the maximum number of translates of P_2 around a vertex, and let k denote the maximum of k_1 and k_2. Then in $\text{star}(P_2)$, we can join any point in the interior of P_2 to any point in the interior of any translate of P_2, by a path which crosses at most k faces. Note that a translate of P_2 in $\text{star}(P_2)$ may have just a vertex in common with P_2. It follows immediately that in $\text{star}^n(P_2)$, we can join P_2 to any translate of P_2, by a path which crosses at most nk faces. Hence, in $\text{star}^n(P_2)$, we can join any two translates of P_2, by a path which crosses at most $2nk$ faces. We conclude that if there is an isomorphism from $\pi_1(M_1)$ to $\pi_1(M_2)$, there is one which maps each face pairing generator of $\pi_1(M_1)$ to a word of length no more than $2nk$ in the face pairing generators of $\pi_1(M_2)$. Similarly we can find integers m and k' such that if there is an isomorphism from $\pi_1(M_2)$ to $\pi_1(M_1)$, there is one which is inverse to the previous one, which maps each face pairing generator of $\pi_1(M_2)$ to a word of length no more than $2mk'$ in the face pairing generators of $\pi_1(M_1)$. This gives us a finite list of possible maps from generators of $\pi_1(M_1)$ to elements of $\pi_1(M_2)$, and vice versa. For each such map we can check whether it is a homomorphism, and for each pair of such maps, can check if their composite is the identity. Each of these checks

again requires the solution of the word problem. Thus we can check whether or not there is an isomorphism between $\pi_1(M_1)$ and $\pi_1(M_2)$.

Before giving the proof of Lemma 2.1, we need to discuss algorithms for bounding various distances related to the convex polyhedra P_i.

We will consider the upper half space model of $\mathbb{H}^3$, and all coordinates used will be euclidean. Recall that in this model a hyperbolic geodesic is either a vertical line, or a semicircle in a vertical plane centred at some point of the base plane $\mathbb{R}^2$ of the model. The two points which form the intersection of such a semicircle with the base plane $\mathbb{R}^2$ of the model will be called the boundary of the geodesic. If a hyperbolic geodesic is a vertical line in this model, its boundary consists of one point in $\mathbb{R}^2$ and one point at infinity. Also a hyperbolic plane in this model is either a vertical plane or is a euclidean hemisphere centred at some point of the base plane $\mathbb{R}^2$ of the model. The circle which is the intersection of such a hemisphere with the base plane $\mathbb{R}^2$ of the model will be called the boundary of the hyperbolic plane. If a hyperbolic plane is a vertical plane in this model, its boundary consists of a line in $\mathbb{R}^2$ and one point at infinity.

We recall that Manning's paper [14] produces the hemispheres which contain the faces of P_i, and that each of these hemispheres has centre with coordinates which are algebraic numbers and has euclidean radius which is also an algebraic number. If the intersection of two of these hemispheres is nonempty, it is a semicircle whose boundary points have coordinates which are algebraic numbers and whose euclidean radius is also an algebraic number. Further the vertices of P_i, each of which is the intersection of three of these hemispheres, also have coordinates which are algebraic numbers. These numbers can be approximated to any required degree of accuracy over the rational numbers (to lie within an "isolating interval" with rational endpoints) using standard methods of symbolic computation, as described by Manning (with reference to Becker and Weispfenning [2] and Loos [13]). The algorithmic process starts from the results of Manning's algorithm, which gives coordinates (ie minimal polynomials and isolating intervals for their roots) for the vertices of the polyhedra P_i, and for the centres of the semicircles and hemispheres defining the edges and faces of P_i, and estimates of their radii. Suppose that all these isolating intervals are of width at most ϵ, and to begin, suppose that ϵ has been chosen so that $1/2^{c+1} < \epsilon < 1/2^c$. We refer to this number as the error in our calculations.

In the upper half-space model, after estimating euclidean distances, we can then estimate hyperbolic distances using the usual formulae. For simplicity in the following seven statements, we will say that a point in the upper half space model of $\mathbb{H}^3$ is algebraic if its coordinates are algebraic, that an infinite geodesic in this model is algebraic if its boundary points are algebraic (we count ∞ as being algebraic here), that a hemisphere

in this model is algebraic if its euclidean centre (in the base plane $z = 0$) and radius are algebraic, and that a vertical plane is algebraic if it contains at least two finite algebraic points (or has at least one finite algebraic boundary point). Finally a compact geodesic segment is algebraic if its endpoints are algebraic.

(1) The distance between two distinct algebraic points is not algorithmically computable, but we can compute a positive lower bound which will suffice for our requirements. The error incurred here is at most 2ϵ.

(2) For this point and the next, we use the upper half space model of the hyperbolic plane $\mathbb{H}^2$. Finding the distance between an algebraic point X and a disjoint algebraic geodesic λ in the hyperbolic plane can be algorithmically reduced to finding the distance between two algebraic points as follows. Let μ denote the semicircle through X which meets λ at right angles and has its centre on the base line $\mathbb{R}$. We can find this centre by solving quadratic equations specifying that the centre is equidistant from X and the point $\lambda \cap \mu$, and that the line from the centre to $\lambda \cap \mu$ is tangent to λ. Hence we can also find $\lambda \cap \mu$. Now the required distance equals the distance from X to the point $\lambda \cap \mu$, so we can apply (1).

(3) In the same way, finding the distance between two disjoint algebraic geodesics (without a common boundary point) in the hyperbolic plane can be algorithmically reduced to (1).

(4) In the same way, finding the distance between an algebraic point X and a disjoint algebraic hyperbolic plane Π in hyperbolic 3–space can be algorithmically reduced to (1).

(5) In the same way, finding the distance between two disjoint algebraic hyperbolic planes whose boundaries are also disjoint can be algorithmically reduced to (1).

(6) To find a lower bound for the distance between an algebraic geodesic λ and a disjoint algebraic hyperbolic plane Π, such that the boundaries of λ and Π are also disjoint, we will choose an algebraic hyperbolic plane Σ which contains λ and is disjoint from Π, so that their boundaries are also disjoint. Once such Σ has been found, the distance between Σ and Π, which can be found as in (5), gives the required lower bound.

If λ is a vertical line, Π must be a hemisphere, and we choose Σ to be the vertical plane through λ which is orthogonal to the vertical plane which contains both λ and the centre of the hemisphere Σ.

Otherwise λ is a semicircle with endpoints a and b in the plane $z = 0$. Now we apply the Moebius transformation φ given by $\varphi(z) = (z-a)/(z-b)$, which takes λ to the vertical line λ' above the origin, and takes Π to an algebraic

hyperbolic plane Π'. The distance between λ' and Π' can be bounded below as in the preceding two paragraphs. As φ is a hyperbolic isometry, this is the required lower bound for the distance between λ and Π.

(7) Finding the distance between an algebraic geodesic segment e and a disjoint algebraic hyperbolic plane Π in hyperbolic 3–space can be algorithmically reduced to the preceding cases as follows. Let λ denote the geodesic which contains e. If λ and Π are disjoint and do not have a common boundary point, we can apply (6) to find a lower bound for the distance between them. This is also a lower bound for the distance between e and Π. Otherwise, the distance between e and Π equals the distance between ∂e and Π, which reduces the problem to (4).

Lemma 2.1 *There is an algorithm to find a positive number R, such that* $\text{star}^n(P_2)$ *contains the metric ball $B(P_2, nR)$.*

Proof We will find R such that for each $n \geq 1$, the R–neighborhood of $\partial \text{star}^n(P_2)$, does not meet $\text{star}^{n-1}(P_2)$. Thus any path in $\mathbb{H}^3$ which starts on $\partial \text{star}^n(P_2)$ and ends on $\partial \text{star}^{n-1}(P_2)$ must have length at least R. By induction it follows that any path which starts on $\partial \text{star}^n(P_2)$ and ends on ∂P_2 must have length at least nR. It follows immediately that $\text{star}^n(P_2)$ contains the metric ball $B(P_2, nR)$, as required.

We first give a description of the exact calculation before considering the error term. We use the above seven points to find positive lower bounds for various distances.

The distance between disjoint vertices of P_2 can be bounded below using (1).

The distance between a vertex v of P_2 and a disjoint edge e of P_2 can be estimated as follows. Let λ denote the geodesic which contains e. As v cannot lie on λ, it suffices to estimate the distance of v from λ. If λ is a vertical line, this can be done as in (2). Otherwise, as in (6), let a and b denote the endpoints of λ, and apply the Moebius transformation φ given by $\varphi(z) = (z-a)/(z-b)$. This takes λ to a vertical line λ' and takes v to an algebraic point v', so we can now estimate the distance of v' from λ' as in (2).

The distance between a vertex v of P_2 and a disjoint face F of P_2 can be estimated using (4), as v cannot lie in the plane which contains F.

The distance between disjoint edges e and f of P_2 can be estimated as follows. Let λ and μ denote the geodesics which contain e and f respectively. If λ meets μ at a finite point or at infinity, the distance between e and f is equal to the distance between ∂e and ∂f, which can be bounded as in (1). If λ and μ are disjoint, and disjoint at infinity, we will find an algebraic plane Π which contains λ and is disjoint

from μ, and is also disjoint from μ at infinity. If λ is a vertical line, so that μ must be a semicircle, we take Π to be the vertical plane through λ which is parallel to the base line of the semicircle μ. If λ is not a vertical line, then, as in (6), we can apply a Moebius transformation which takes λ to a vertical line and takes μ to an algebraic geodesic.

The distance between an edge e of P_2 and a disjoint face F of P_2 can be estimated using (7), as e must be disjoint from the plane which contains F.

Finally the distance between disjoint faces E and F of P_2 can be estimated as follows. Let Π_E and Π_F denote the planes which contain E and F respectively. If these planes are disjoint, and disjoint at infinity, we can estimate the distance between them using (5), and this will be a lower bound for the distance between E and F. Otherwise, the distance between E and F is bounded below by the distance between E and Π_F. This last distance equals the distance between Π_F and some edge of E, and so equals one of the numbers already estimated.

Now let R denote half the minimum of all these numbers.

Let W denote a vertex, edge or face of P_2. Then the definition of R implies that the R–neighborhood of W meets only those vertices, edges or faces of P_2 which meet W.

Hence if Q is a translate of P_2 in $\mathrm{star}^n(P_2)$ which meets $\partial\,\mathrm{star}^n(P_2)$, then the R–neighborhood of $\partial\,\mathrm{star}^n(P_2)$ does not meet any vertex, edge or face of Q except those which meet $\partial\,\mathrm{star}^n(P_2)$. In particular, it follows that the R–neighborhood of $\partial\,\mathrm{star}^n(P_2)$, does not meet $\mathrm{star}^{n-1}(P_2)$, as required. Note that both $\partial\,\mathrm{star}^n(P_2)$ and $\mathrm{star}^{n-1}(P_2)$ are disjoint.

In the actual algorithm, when dealing with approximations, all the calculations above incur increasing error, but the fact that the number of operations is finite means that there is a constant $C > 0$ such that the error in the estimate of each of these numbers is at most $C\epsilon$, so that R must be replaced by $R - C\epsilon$. It is of course possible that $R < C\epsilon$ in which case the algorithm must restart, replacing ϵ by $\epsilon/2$, recalculating the coordinates in Manning's algorithm to this increased degree of accuracy, and recalculating R. Continue to do so until $R > C\epsilon$, and then replace R by $R - C\epsilon$, once this number is positive. □

3   Algorithm to find the geometry

In this section we consider closed orientable irreducible geometric 3–manifolds, given by finite triangulations. There are eight geometries (as discussed in the first author's

article [21]), and it is well known that a closed 3–manifold can have a geometric structure modelled on at most one of these geometries. This can be proved by exhibiting properties of the fundamental groups which distinguish the geometries. For example only closed manifolds modelled on S^3 can have finite fundamental group. The point of what we do in this section is that we can decide algorithmically on which geometry a given geometric manifold is modelled. Such an algorithm is described by Manning in [14, 1.4], but here we provide a more algebraic treatment.

As we are considering orientable irreducible 3–manifolds, the geometry $S^2 \times \mathbb{R}$ cannot occur. For the only closed orientable manifolds modelled on this geometry are $S^2 \times S^1$ and $\mathbb{R}P^3 \# \mathbb{R}P^3$, neither of which is irreducible. We start by listing the remaining seven geometries together with some selected properties of the closed manifolds modelled on these geometries.

Geometry	Selected properties of any closed orientable 3–manifold M modelled on given geometry
S^3	$\pi_1(M)$ is finite.
$\mathbb{E}^3$	$\pi_1(M)$ is virtually $\mathbb{Z}^3$ and M is Haken. Any two–sided incompressible surface in M must be a torus.
$\mathbb{H}^3$	$\pi_1(M)$ has no subgroup isomorphic to $\mathbb{Z}^2$.
$\mathbb{H}^2 \times \mathbb{R}$	M is a Seifert fibre space with hyperbolic base orbifold. M is Haken and contains an embedded incompressible hyperbolic surface.
Nil	M is a Seifert fibre space with euclidean base orbifold.
Solv	M is Haken. Any two–sided incompressible surface in M must be a torus.
$\widetilde{SL_2\mathbb{R}}$	M is a Seifert fibre space with hyperbolic base orbifold. A two–sided incompressible surface in M must be a torus.

A crucial fact for us is that if M is modelled on one of the above seven geometries, then $\pi_1(M)$ is biautomatic, except in the cases when the geometry is Nil or Solv ([6, Chapter 12] proves automaticity). In order to apply the theory of biautomatic structures, we first need to be able to decide whether M is modelled on Nil or Solv.

The key topological algorithm we will need is that of Jaco and Oertel [11] which decides whether a given triangulated 3–manifold M is Haken. Their paper also shows how to decide whether M has an incompressible surface which is a torus and how to find such a torus. We will use these algorithms several times in what follows.

As usual T denotes the 2–torus $S^1 \times S^1$. We also need to be able to decide whether a compact orientable manifold M' is homeomorphic to $T \times I$. If M' is irreducible, this is a special case of [12, Algorithm 9.7]. Note that the algorithms of Jaco and Oertel also find essential discs, and in this case cutting along a properly embedded disc gives a 3–ball, which can be recognised (by Rubinstein [20] and Thompson [23], or by Perelman's solution of the Poincaré conjecture).

Lemma 3.1 [21, Theorem 5.5] *If M is orientable and is obtained from $T \times I$ by gluing $T \times \{0\}$ to $T \times \{1\}$ by some homeomorphism h, then M is geometric and is modelled on one of $\mathbb{E}^3$, Nil or Solv.*

Proof The action of h on $H_1(T) \cong \mathbb{Z}^2$ is given by an integer 2×2 matrix A. We consider the trace, $\mathrm{tr}(A)$, of A. If $|\mathrm{tr}(A)| < 2$, or if $A = \pm I$, then A, and hence h, must be periodic, so that M is modelled on $\mathbb{E}^3$. If $|\mathrm{tr}(A)| > 2$, then A has distinct real eigenvalues, so that M is modelled on Solv. If $|\mathrm{tr}(A)| = 2$, then A has a repeated eigenvalue equal to ± 1. So long as $A \neq \pm I$, this implies M is modelled on Nil. $\square$

Remark Suppose that we have found an incompressible torus T in an orientable 3–manifold M, using the algorithms of normal surface theory, and that we have checked that cutting M along T yields a manifold homeomorphic to $T \times I$. In this situation, one can algorithmically calculate the action of h on $H_1(T) \cong \mathbb{Z}^2$, and so can decide on which geometry M is modelled.

We now discuss the geometries Nil and Solv in more detail and describe an algorithm to decide which geometry occurs.

If M is modelled on Solv, then M is a bundle over a 1–dimensional orbifold with fibre the torus. Thus either M is a bundle over S^1 with fibre the torus, or M is double covered by such a manifold.

If M is modelled on Nil, there are several cases. If M is Haken, then it is a Seifert fibre space whose base orbifold is a torus, Klein bottle, $S^2(2, 2, 2, 2)$ or $P^2(2, 2)$. If this orbifold is not a torus, there is a regular cover of M of degree 2 or 4 whose base orbifold is a torus. If M is not Haken, it is a Seifert fibre space whose base orbifold is $S^2(p, q, r)$, where (p, q, r) is one of $(3, 3, 3)$, $(2, 2, 4)$ or $(2, 3, 6)$. Now the orbifold fundamental group of $S^2(p, q, r)$ is the triangle group $\Delta(p, q, r)$, and in these cases, $\Delta(p, q, r)$ has a homomorphism to $\mathbb{Z}_3$, $\mathbb{Z}_4$ or $\mathbb{Z}_6$ with kernel isomorphic to $\mathbb{Z}^2$. Thus there is a homomorphism of $\pi_1(M)$ to $\{1\}$, $\mathbb{Z}_3$, $\mathbb{Z}_4$ or $\mathbb{Z}_6$ whose kernel determines a finite cover of M which is a circle bundle over the torus. If M is modelled on Nil and is a circle bundle over the torus, then any two–sided incompressible surface

in M must be a vertical torus, and if we cut M along such a torus, the result will be homeomorphic to $T \times I$.

We can apply the preceding paragraph as follows. For any triangulated closed orientable irreducible 3–manifold M, we can check whether there is a homomorphism of $\pi_1(M)$ to $\{1\}$, $\mathbb{Z}_2$, $\mathbb{Z}_3$, $\mathbb{Z}_4$, $\mathbb{Z}_2 \times \mathbb{Z}_2$, or $\mathbb{Z}_6$ whose kernel determines a finite cover of M with infinite first homology group. If this does not occur, then M cannot be modelled on Nil or Solv. If this does occur, we can check whether the covering contains an incompressible torus, and if it does, we can check whether cutting along this torus yields $T \times I$. If this occurs, the remark above tells us how to determine the geometry on this finite cover and hence on M.

Thus we can decide whether or not M is modelled on Nil or Solv, and if it is so modelled can decide which.

This reduces us to considering the five remaining geometries.

Given a triangulated closed orientable irreducible 3–manifold M, we can write down a presentation for $\pi_1(M)$. If we know that $\pi_1(M)$ is biautomatic, we can algorithmically find a biautomatic structure. Part of this structure is a regular (or rational) language of representatives for the elements of the group. We can suppose, using [6, Theorem 2.5.2], that each group element has a unique representative in the language, and it is easy to check whether a regular language is finite or infinite (see for instance Hopcroft and Ullman [9, Theorem 3.7]). Thus we can algorithmically check whether $\pi_1(M)$ is finite, and so can decide whether M is modelled on S^3.

This reduces us to the four remaining geometries, which are $\mathbb{E}^3$, $\mathbb{H}^3$, $\mathbb{H}^2 \times \mathbb{R}$ and $\widetilde{\mathrm{SL}_2\mathbb{R}}$.

If M is modelled on $\mathbb{E}^3$, there are several cases. In all cases, M is Haken. It is a Seifert fibre space whose base orbifold is a torus, Klein bottle, $S^2(2,2,2,2)$, $P^2(2,2)$, or $S^2(p,q,r)$, where (p,q,r) is one of $(3,3,3)$, $(2,2,4)$ or $(2,3,6)$. Thus, as for Nil geometry, there is a homomorphism of $\pi_1(M)$ to $\{1\}$, $\mathbb{Z}_2$, $\mathbb{Z}_3$, $\mathbb{Z}_4$, $\mathbb{Z}_2 \times \mathbb{Z}_2$, or $\mathbb{Z}_6$ whose kernel determines a finite cover of M whose base orbifold is the torus. But now this finite cover must be a 3–torus, and so have free abelian fundamental group.

Thus to decide whether or not M is modelled on $\mathbb{E}^3$, we simply check whether there is a homomorphism of $\pi_1(M)$ to $\mathbb{Z}_2$, $\mathbb{Z}_3$, $\mathbb{Z}_4$, $\mathbb{Z}_2 \times \mathbb{Z}_2$, or $\mathbb{Z}_6$ whose kernel is free abelian of rank 3. For M is modelled on $\mathbb{E}^3$, if and only if there is such a homomorphism.

This reduces us to the three remaining geometries, which are $\mathbb{H}^3$, $\mathbb{H}^2 \times \mathbb{R}$ and $\widetilde{\mathrm{SL}_2\mathbb{R}}$.

Using the biautomatic structure on $\pi_1(M)$ one can check the answers to the following questions. Does $\pi_1(M)$ have nontrivial centre? Does $\pi_1(M)$ have a subgroup of index 2 with nontrivial centre? (A regular language for the centre of a biautomatic group is constructed in Gersten and the second author [7, Corollary 4.4.1], and as noted earlier, it is easy to check whether the language is infinite or finite, and in the latter case deduce how many elements are in the centre. It is also straightforward to obtain a biautomatic structure for all subgroups of a given finite index, as in [6, Theorem 4.1.1].) If the answer to both questions is negative, then M must be hyperbolic. If we find a positive answer, then M must be modelled on one of $\mathbb{H}^2 \times \mathbb{R}$ and $\widetilde{\mathrm{SL}_2\mathbb{R}}$. To distinguish these cases, we use the facts that if M is modelled on $\widetilde{\mathrm{SL}_2\mathbb{R}}$, then any incompressible surface in M must be a (vertical) torus, whereas if M is modelled on $\mathbb{H}^2 \times \mathbb{R}$, there must be horizontal incompressible surfaces in M none of which can be a torus. Thus we can apply the algorithm of Jaco and Oertel [11] to decide whether M contains an incompressible surface which is not a torus.

The referee pointed out an alternative algebraic approach to distinguishing the $\mathbb{H}^2 \times \mathbb{R}$ and $\widetilde{\mathrm{SL}_2\mathbb{R}}$ cases. It is based on two observations. The first is that if M is modelled on one of these two geometries, then M has a finite cover M_1 which is a bundle over a surface with fibre the circle, such that the centre of $\pi_1(M_1)$ is infinite cyclic. The second is that if M_1 is such a manifold, one can decide on which geometry M_1 (and hence M) is modelled by checking whether the centre of $\pi_1(M_1)$ injects into $H_1(M_1)$. If it does, then the geometry is $\mathbb{H}^2 \times \mathbb{R}$, and if it does not, then the geometry is $\widetilde{\mathrm{SL}_2\mathbb{R}}$. It will be helpful to add a third observation. This is that if the centre of $\pi_1(M)$ is infinite cyclic, and if the centre of $\pi_1(M)$ injects into $H_1(M)$, then M must be modelled on $\mathbb{H}^2 \times \mathbb{R}$. This is true because the assumption that the centre of $\pi_1(M)$ injects into $H_1(M)$ immediately implies that the centre of $\pi_1(M_1)$ injects into $H_1(M_1)$.

To distinguish these two geometries algorithmically, we proceed as follows. Suppose that $\pi_1(M)$ has nontrivial centre. (If not, replace M by the double cover which does have this property.) Now this centre, which we denote by A, is infinite cyclic and the quotient $\pi_1(M)/A$ is a Fuchsian group Γ. If Γ is torsion free, it is a surface group, and we take M_1 to equal M. We can decide whether Γ is torsion free. First Γ is δ–hyperbolic for some $\delta > 0$, and such a δ can be found algorithmically (see for instance [16]). Now any torsion element of Γ must have length at most $4\delta + 2$ (Bridson and Haefliger [4, proof of Theorem III.Γ.3.2]) and so the orders of all torsion elements can be found. If there is any nontrivial torsion, we let k denote the least common multiple of these orders. It follows from Edmonds, Ewing and Kulkarni [5, Theorem 1] that there is a torsion free subgroup Γ_1 of index $2k$ in Γ. Thus there is a degree $2k$ cover M_1 of M, such that $\pi_1(M_1)$ has centre A, and $\pi_1(M_1)/A$ is the surface group Γ_1. We can now algorithmically find finite presentations for the (finitely many)

subgroups of Γ of index $2k$, one of which is Γ_1. At this point one could simply check each of these subgroups of Γ to decide whether it is torsion free, but it seems simpler to proceed in the following way. Instead consider the index $2k$ subgroups of $\pi_1(M)$ with centre A and quotient one of the index $2k$ subgroups of Γ. For each such subgroup G of $\pi_1(M)$, we check whether A injects into the abelianisation of G. If this does not occur for any of these subgroups we immediately deduce that M has $\widetilde{\mathrm{SL}_2\mathbb{R}}$ geometry. If this does occur for some such subgroup, say $\pi_1(M_2)$, we can apply the third observation above to see that M_2, and hence M, has $\mathbb{H}^2 \times \mathbb{R}$ geometry, without ever needing to check whether the quotient $\pi_1(M_2)/A$ is torsion free.

References

[1] **J M Alonso**, **T Brady**, **D Cooper**, **V Ferlini**, **M Lustig**, **M Mihalik**, **M Shapiro**, **H Short**, *Notes on word hyperbolic groups*, from: "Group theory from a geometrical viewpoint", (H Short, editor), World Sci. Publ. (1991) 3–63   MR1170363

[2] **T Becker**, **V Weispfenning**, *Gröbner bases: A computational approach to commutative algebra*, Graduate Texts in Math. 141, Springer, New York (1993)   MR1213453

[3] **L Bessières**, **G Besson**, **S Maillot**, **M Boileau**, **J Porti**, *Geometrisation of* 3*–manifolds*, EMS Tracts in Math. 13, European Math. Soc. (EMS), Zürich (2010)   MR2683385

[4] **M R Bridson**, **A Haefliger**, *Metric spaces of non-positive curvature*, Grundl. Math. Wissen. 319, Springer, Berlin (1999)   MR1744486

[5] **A L Edmonds**, **J H Ewing**, **R S Kulkarni**, *Torsion free subgroups of Fuchsian groups and tessellations of surfaces*, Bull. Amer. Math. Soc. 6 (1982) 456–458   MR648534

[6] **D B A Epstein**, **J W Cannon**, **D F Holt**, **S V F Levy**, **M S Paterson**, **W P Thurston**, *Word processing in groups*, Jones and Bartlett Publishers, Boston (1992)   MR1161694

[7] **S M Gersten**, **H B Short**, *Rational subgroups of biautomatic groups*, Ann. of Math. 134 (1991) 125–158   MR1114609

[8] **S M Gersten**, **H Short**, *Small cancellation theory and automatic groups, II*, Invent. Math. 105 (1991) 641–662   MR1117155

[9] **J E Hopcroft**, **J D Ullman**, *Introduction to automata theory, languages, and computation*, Addison-Wesley Publishing Co., Reading, MA (1979)   MR645539

[10] **W Jaco**, *Peking summer school lectures*, preprint (2005)   Available at http://www.math.okstate.edu/~jaco/pekinglectures.htm

[11] **W Jaco**, **U Oertel**, *An algorithm to decide if a* 3*–manifold is a Haken manifold*, Topology 23 (1984) 195–209   MR744850

[12] **W Jaco**, **J L Tollefson**, *Algorithms for the complete decomposition of a closed* 3*–manifold*, Illinois J. Math. 39 (1995) 358–406   MR1339832

[13] **R Loos**, *Computing in algebraic extensions*, from: "Computer algebra", (B Buchberger, G E Collins, R Loos, R Albrecht, editors), Springer, Vienna (1983) 173–187 MR728972

[14] **J Manning**, *Algorithmic detection and description of hyperbolic structures on closed 3–manifolds with solvable word problem*, Geom. Topol. 6 (2002) 1–25   MR1885587

[15] **S Matveev**, *Algorithmic topology and classification of 3-manifolds*, 2nd edition, Algorithms and Computation in Math. 9, Springer, Berlin (2007)   MR2341532

[16] **P Papasoglu**, *An algorithm detecting hyperbolicity*, from: "Geometric and computational perspectives on infinite groups", (G Baumslag, D Epstein, R Gilman, H Short, C Sims, editors), DIMACS Ser. Discrete Math. Theoret. Comput. Sci. 25, Amer. Math. Soc. (1996) 193–200   MR1364185

[17] **G Perelman**, *The entropy formula for the Ricci flow and its geometric applications* (2002)   arXiv:math/0211159

[18] **G Perelman**, *Finite extinction time for the solutions to the Ricci flow on certain* 3–*manifolds* (2002)   arXiv:math/0307245

[19] **G Perelman**, *Ricci flow with surgery on* 3–*manifolds* (2002)   arXiv:math/0303109

[20] **J H Rubinstein**, *An algorithm to recognize the* 3–*sphere*, from: "Proceedings of the International Congress of Mathematicians, Vol. 1, 2", (S D Chatterji, editor), Birkhäuser, Basel (1995) 601–611   MR1403961

[21] **P Scott**, *The geometries of* 3–*manifolds*, Bull. London Math. Soc. 15 (1983) 401–487 MR705527

[22] **Z Sela**, *The isomorphism problem for hyperbolic groups, I*, Ann. of Math. 141 (1995) 217–283   MR1324134

[23] **A Thompson**, *Thin position and the recognition problem for* S^3, Math. Res. Lett. 1 (1994) 613–630   MR1295555

Mathematics Department, University of Michigan at Ann Arbor,
Ann Arbor, MI 48109, USA

CMI, Universite d'aix-Marseille,
UMR 7353, 39 Rue Joliot Curie, 13453 Marseille, France

pscott@umich.edu,   hamish.short@univ-amu.fr