

Recognising elliptic manifolds

Marc Lackenby and Saul Schleimer

Abstract. We show that the problem of deciding whether a closed three-manifold admits an elliptic structure lies in NP. Furthermore, determining the homeomorphism type of an elliptic manifold lies in the complexity class FNP. These are both consequences of the following result. Suppose that M is a lens space which is neither $\mathbb{R}P^3$ nor a prism manifold. Suppose that \mathcal{T} is a triangulation of M . Then there is a loop, in the one-skeleton of the 86th iterated barycentric subdivision of \mathcal{T} , whose simplicial neighbourhood is a Heegaard solid torus for M .

1. Introduction

Compact orientable three-manifolds have been classified in the following sense: there are algorithms that, given two such manifolds, determine if they are homeomorphic [1, 13, 27]. Kuperberg [13, Theorem 1.2] has further shown that this problem is no worse than elementary recursive. Beyond this, very little is known about the computational complexity of the homeomorphism problem.

All known solutions rely on the geometrization theorem, due to Perelman [21–23]. This motivates the following closely related problem: given a compact orientable three-manifold, determine if it admits one of the eight Thurston geometries. This problem has an exponential-time solution using normal surface theory and, again, geometrization. See [13, Section 8.2] for a closely related discussion.

Here we give a much better upper bound in an important special case. Recall that the *elliptic* manifolds are those admitting spherical geometry. The decision problem ELLIPTIC MANIFOLD takes as input a triangulation \mathcal{T} , of a compact connected three-manifold M , and asks if M is elliptic.

Theorem 1.1. *The problem ELLIPTIC MANIFOLD lies in NP.*

If a three-manifold is elliptic, then it is reasonable to ask which elliptic manifold it is. This is a *function problem* as the desired output is more complicated than simply

“yes” or “no”. The problem NAMING ELLIPTIC takes as input a triangulation of a compact connected three-manifold, which is promised to admit an elliptic structure, and requires as output the manifold’s Seifert data. Some elliptic three-manifolds admit more than one Seifert fibration; in this case, the output is permitted to be the data for any of these.

Theorem 1.2. *The problem NAMING ELLIPTIC lies in FNP.*

One precursor to Theorem 1.1 is that THREE-SPHERE RECOGNITION lies in NP. The first proof of this [25, Theorem 15.1] uses Casson’s technique of *crushing* normal two-spheres as well as Rubinstein’s *sweep-outs*, derived from almost normal two-spheres. There is another proof, due to Ivanov [9, Theorem 2], that again uses crushing but avoids the machinery of sweep-outs. Ivanov also shows that the problem of recognising the solid torus lies in NP.

Our results rely on this prior work in a crucial but non-obvious fashion. By geometrisation, a three-manifold M is elliptic if and only if M is finitely covered by the three-sphere. Thus, one might hope to prove Theorems 1.1 and 1.2 by exhibiting such a finite cover together with a certificate that the cover is S^3 . However, consider the following examples. Let F_n denote the n th Fibonacci number. There is a triangulation of the lens space $L(F_n, F_{n-1})$ with n tetrahedra. The degree of the universal covering is F_n ; since this grows exponentially in n it cannot be used directly in an NP certificate. See [12, Section 2] for many more examples of this phenomenon.

Instead we use the following: any elliptic three-manifold has a cover, of degree at most sixty, which is a lens space. Thus Theorems 1.1 and 1.2 reduce, respectively, to the problem of deciding whether a three-manifold is a lens space and, if so, naming it.

Our approach to certifying lens spaces is conceptually simple. Suppose that

$$U \cong V \cong S^1 \times D^2$$

are solid tori. A lens space M can be obtained by gluing U and V along their boundaries. This gives a *Heegaard splitting* for M . This decomposition of M is unique up to ambient isotopy [2, Théorème 1]. We call

$$S^1 \times \{0\} \subset S^1 \times D^2$$

a *core curve* for the solid torus. We say that a simple closed curve $\gamma \subset M$ is a *core curve* for M if it is isotopic to a core curve for U or for V . We certify that M is a lens space by exhibiting such a core curve.

This approach is inspired by the results of [14]. There Lackenby shows that, for any handle structure of a solid torus satisfying some natural conditions, there is a core curve that lies nicely with respect to the handles. Specifically, the curve lies within the union of the zero- and one-handles and its intersection with each such handle is one of

finitely many types. This list of types is universal, in the sense that it does not depend on the handle structure. The handle structure must satisfy some hypotheses, but these hold for any handle structure that is dual to a triangulation. Using [14, Theorem 4.2] we give an explicit bound on the “combinatorial length” of the core curve. For a triangulation \mathcal{T} and positive integer n , we let $\mathcal{T}^{(n)}$ denote the triangulation obtained from \mathcal{T} by performing barycentric subdivision n times.

Theorem 1.3. *Suppose that \mathcal{T} is a triangulation of the solid torus M . Then M contains a core curve that is a subcomplex of $\mathcal{T}^{(51)}$.*

Using this we prove the following technical result.

Theorem 1.4. *Suppose that M is a lens space other than $\mathbb{R}\mathbb{P}^3$. Suppose that \mathcal{T} is a triangulation of M . Then there is a simple closed curve C that is a subcomplex of $\mathcal{T}^{(86)}$, such that the exterior of C is either a solid torus or a twisted I -bundle over a Klein bottle.*

This is proved by placing a Heegaard torus S into almost normal form in the triangulation \mathcal{T} , using Stocking’s work [28, Theorem 1]. We then cut along this torus to get two solid tori, which inherit handle structures \mathcal{H}_1 and \mathcal{H}_2 . We could apply Theorem 1.3 to each of these, and we would then obtain core curves of the solid tori. However, their intersection with each tetrahedron of \mathcal{T} would not necessarily be of the required form. In particular, the intersection between each of these curves and any tetrahedron of \mathcal{T} would not be in one of finitely many configurations; thus there would be no way of showing that it lay in $\mathcal{T}^{(86)}$. The reason for this is that each tetrahedron of \mathcal{T} may contain many handles of \mathcal{H}_1 and \mathcal{H}_2 . However, in this situation, all but a bounded number of these handles would lie between normally parallel triangles and squares of S . Hence, they lie in the *parallelity bundle* for \mathcal{H}_1 or \mathcal{H}_2 . (The parallelity bundle was first introduced by Kneser [11]; we follow the treatment given in [15].) The strategy behind Theorem 1.4 is to ensure that one of the core curves does not intersect these handles by using the results from [14]. It then intersects each tetrahedron of \mathcal{T} in one of finitely many types, and with some work, we show that it is in fact simplicial in $\mathcal{T}^{(86)}$.

Using Theorems 1.3 and 1.4 we prove our main technical result.

Theorem 1.5. *Suppose that M is a lens space which is neither a prism manifold nor a copy of $\mathbb{R}\mathbb{P}^3$. Suppose that \mathcal{T} is a triangulation of M . Then the iterated barycentric subdivision $\mathcal{T}^{(86)}$ contains a core curve of M in its one-skeleton. Furthermore, $\mathcal{T}^{(139)}$ contains in its one-skeleton the union of the two core curves.*

1.1. Other work

We announced Theorems 1.1 and 1.2 in 2012, at Oberwolfach [17]. Motivated by this, Kuperberg [12, Theorem 1.1] showed that the *function promise problem* NAMING LENS SPACE has a polynomial-time solution. That is, there is a polynomial-time algorithm that, given a triangulated lens space M , determines its lens space coefficients. His work, together with Theorem 1.1 and parts of Section 12, can be used to give another proof of Theorem 1.2.

The work of Haraway and Hoffman [8] is also relevant here. In particular, in Section 11, we rely upon their result [8, Theorem 3.6] that the decision problems RECOGNISING $T^2 \times I$ and RECOGNISING $K^2 \tilde{\times} I$ lie in NP.

1.2. Outline of paper

We assume that all of our three-manifolds are given via a triangulation; we work in the piecewise-linear category throughout.

In Section 2, we remind the reader of some elementary facts about lens spaces and *prism manifolds*. The lens spaces which are also prism manifolds are exceptional cases in our analysis (as can be seen in the statement of Theorem 1.5, for example). In Section 3, we recall the definition of an (almost) normal surface in a triangulated three-manifold. In Section 4, we discuss handle structures for three-manifolds. In Section 5, we give the background needed to state [14, Theorem 4.2]. This result places the core curve of a solid torus into a controlled position with respect to a handle structure for the solid torus. In Section 6, we translate this back to triangulations. In particular, we prove Theorem 1.3.

In Section 7, we say what it means for one three-manifold to be *nicely embedded* into another. Our eventual goal is to show that one of the Heegaard solid tori of the lens space M is nicely embedded within the triangulation \mathcal{T} of M . We show that, in this situation, the core curve of this solid torus can be arranged to be simplicial in an iterated barycentric subdivision of \mathcal{T} .

In Section 8, we recall the notion of parallelity bundles in a handle structure. In addition, we also discuss *generalised* parallelity bundles, also from [15], where it was shown that such bundles may be chosen to have incompressible horizontal boundary [15, Proposition 5.6]. In our situation, this implies that the generalised parallelity bundle has horizontal boundary being a collection of annuli and discs lying in the Heegaard torus. In Section 9, we bring these results together and prove Theorems 1.5 and 1.4. This section may be viewed as the heart of the paper.

In Section 10, we recall the classification of elliptic three-manifolds into lens spaces, prism manifolds, and *platonic* manifolds. We show that any platonic manifold has a cover, of degree at most sixty, which is a lens space. We also show how

its Seifert data can be extracted from the cover and from the homology of the manifold. In Section 12, we certify elliptic manifolds, thereby completing the proof of Theorems 1.1 and 1.2.

1.3. The role of geometrisation in this paper

Very little of this paper relies on the proof of Thurston’s geometrisation conjecture by Perelman [21–23]. The proofs of Theorems 1.3, 1.4 and 1.5 make no appeal to geometrisation. As a result, our proof that lens space recognition is in NP also does not rely on geometrisation. As explained above, we certify other elliptic manifolds by exhibiting a cover that is a lens space. The fact that any manifold covered by a lens space is elliptic seems to rely on the resolution of the spherical space form conjecture, which was a consequence of geometrisation. However, the certification of prism manifolds can avoid the use of geometrisation by appealing to earlier results on the spherical space form conjecture by Livesay [18] and Myers [19]. However, the certification of platonic manifolds *does* seem to rely on geometrisation; see the proof of Proposition 10.9.

2. Lens spaces and prism manifolds

In this section, we gather a few facts about lens spaces and prism manifolds.

Fix an orientation of the two-torus T . Suppose that λ and μ are simple closed oriented curves in T . We write $\lambda \cdot \mu$ for the algebraic intersection number of λ and μ . If $\lambda \cdot \mu = 1$, then we call the ordered pair $\langle \lambda, \mu \rangle$ a *framing* of T . In this case, we may isotope λ and μ , so that $x = \lambda \cap \mu$ is a single point. With this done, λ and μ generate

$$\pi_1(T, x) \cong \mathbb{Z}^2.$$

Thus, for any simple closed oriented essential curve α in T , we may write

$$\alpha = p\lambda + q\mu,$$

with p and q coprime. Note that $\lambda \cdot \alpha = q$ and $\alpha \cdot \mu = p$. We say that α has *slope* q/p . Note that if β has slope s/r , then $\alpha \cdot \beta = \pm(ps - qr)$. We say that α and β are *Farey neighbours* if $\alpha \cdot \beta = \pm 1$.

Suppose $U = D^2 \times S^1$ is a solid torus. Fix $x \in \partial D^2$ and $y \in S^1$. The boundary $T = \partial U$ has a framing coming from taking $\lambda = \{x\} \times S^1$ and $\mu = \partial D^2 \times \{y\}$.

Definition 2.1. A three-manifold M , obtained by gluing a pair of solid tori U and V via a homeomorphism of their boundaries, is called a *lens space* if it has finite fundamental group. ■

Under this definition $S^2 \times S^1$ is not a lens space; we use this convention because we are here interested in elliptic manifolds. We use $L(p, q)$ to denote the lens space obtained by attaching the meridian slope in ∂V to the slope q/p in ∂U . We call p and q *coefficients* for the lens space. We now record several facts.

Fact 2.2. We have $\pi_1(L(p, q)) \cong \mathbb{Z}/p\mathbb{Z}$. ■

Note for any q , we have $L(1, q) \cong S^3$.

Fact 2.3 ([24]). The lens space $L(p', q')$ is homeomorphic to $L(p, q)$ if and only if $|p'| = |p|$ and $q' = \pm q^{\pm 1} \bmod p$. ■

Fact 2.4. The double cover of $L(2p, q)$ is homeomorphic to $L(p, q)$. ■

Notice $L(p, q)$ double covers both $L(2p, q)$ and $L(2p, p - q)$. For example, $L(8, 1)$ and $L(8, 3)$ are both covered by $L(4, 1)$. However, $L(8, 1)$ and $L(8, 3)$ are not homeomorphic according to Fact 2.3. Thus one cannot recover the coefficients of a lens space just by knowing a double cover.

Fact 2.5. Suppose that α and β are Farey neighbours in T , with slopes q/p and s/r . Suppose that γ has slope q'/p' in T . The three-manifold $M = U \cup (T \times I) \cup V$, formed by attaching the meridian of U along $\alpha \times 0$ and attaching the meridian of V along $\gamma \times 1$, is homeomorphic to the lens space $L(-p'q + q'p, p's - q'r)$. ■

We use $\tilde{\times}$ to denote a twisted product. We write $K = K^2 = S^1 \tilde{\times} S^1$ for the Klein bottle. We write $K \tilde{\times} I$ for the orientation I -bundle over the Klein bottle.

Recall that K contains exactly four essential simple closed curves, up to isotopy. These are the cores α and α' of the two Möbius bands, their common boundary δ , and the fibre β of the bundle structure $K = S^1 \tilde{\times} S^1$. Thus $\pi_1(K) \cong \pi_1(K \tilde{\times} I)$ has a presentation

$$\langle a, b \mid aba^{-1} = b^{-1} \rangle$$

where $a = [\alpha]$ and $b = [\beta]$. This presentation is not canonical, as we could have chosen α' instead of α . Let $\rho: T \rightarrow K$ be the orientation double cover. Thus we have

$$\rho_*(\pi_1(T)) \cong \langle a^2, b \rangle < \pi_1(K).$$

Since $a^2 = [\delta]$ this generating set for $\pi_1(T)$ gives a canonical framing of T (up to the choice of orientations). The identification of T and $\partial(K \tilde{\times} I)$ now gives us a canonical framing of the latter.

Definition 2.6. A three-manifold M , obtained by gluing an I -bundle $K \tilde{\times} I$ to a solid torus W via a homeomorphism of their boundaries, is called a *prism manifold* if it has finite fundamental group. ■

We use the notation $P(p, q)$ to denote the three-manifold obtained by gluing the meridian slope of a solid torus W to the slope $a^{2p}b^q$ in $\partial(K \tilde{\times} I)$. We call p and q *coefficients* of $P(p, q)$.

Fact 2.7. We have $\pi_1(P(p, q)) \cong \langle a, b \mid aba^{-1} = b^{-1}, a^{2p} = b^{-q} \rangle$. ■

For example, we have $P(1, 1) \cong L(4, 3)$; see Lemma 2.9. On the other hand we have $P(1, 0) \cong \mathbb{RP}^3 \# \mathbb{RP}^3$ and $P(0, 1) \cong S^2 \times S^1$. Since the fundamental groups are infinite, we do not admit $P(1, 0)$ or $P(0, 1)$ as prism manifolds.

Lemma 2.8. *Suppose that M is an irreducible oriented closed non-Haken three-manifold. Then M contains an embedded Klein bottle if and only if it is a prism manifold. In particular, a lens space contains an embedded Klein bottle if and only if it is a prism manifold.*

Proof. By construction, prism manifolds contain an embedded Klein bottle. We need to establish the converse. Suppose that M contains an embedded Klein bottle K . By the orientability of M , the regular neighbourhood $N(K)$ is the orientable I -bundle over the Klein bottle. Let $T = \partial N(K)$ be the boundary torus. Since M is non-Haken, T is compressible along a disc D . Note that D has interior disjoint from $N(K)$. Compressing T along D gives a sphere, which bounds a ball B by the irreducibility of M . Note that B is disjoint from K . Thus T bounds a solid torus with interior disjoint from $N(K)$. Therefore, M is a prism manifold. ■

We now give a proof of [3, Corollary 6.4]. For another account, with further references, see [5, Theorem 1.2]. This proof is included for the convenience for the reader.

Lemma 2.9. *The three-manifold $P(p, q)$ is a lens space if and only if $q = 1$ and $p \neq 0$. In this case, $P(p, 1) \cong L(4p, 2p + 1)$.*

Proof. Recall $\gcd(p, q) = 1$ for any slope q/p . So if $q = 0$, then $p = 1$; likewise, if $p = 0$, then $q = 1$. In these cases, as noted above, $P(1, 0)$ and $P(0, 1)$ are not lens spaces.

Suppose that $q \geq 2$. Taking $P = P(p, q)$ we note that $\pi_1(P)$ has as a quotient group

$$D_{2q} = \langle a, b \mid aba^{-1} = b^{-1}, a^2 = b^q = 1 \rangle$$

the dihedral group of order $2q$. This is not cyclic, so P is not a lens space.

Thus we may assume that $q = 1$ and $p \neq 0$. Set $P = P(p, 1)$. We note that

$$\pi_1(P) = \langle a, b \mid aba^{-1} = b^{-1}, a^{2p} = b^{-1} \rangle \cong \langle a \mid a^{4p} = 1 \rangle.$$

This implies that $P(p, 1)$ is the unique manifold, up to homeomorphism, that is a lens space and prism manifold with fundamental group of order $4p$. We now give a

direct proof that $P = P(p, 1)$ is a lens space. Recall that if we glue a pair of solid tori along an annulus, primitive in at least one of them, the result is a solid torus. Write $K \tilde{\times} I = Q \cup R$ as a union of solid tori, each the orientation I -bundle over a Möbius band. Note $Q \cap R$ is a vertical annulus in $K \tilde{\times} I$. Set $A = Q \cap \partial(K \tilde{\times} I)$. Note that A is not primitive in Q , as it crosses the meridian disc of Q twice.

Recall that $P = (Q \cup R) \cup W$, all solid tori. When $q = 1$, we find that the annulus A is primitive in W . Thus $Q \cup W$ is a solid torus and so $P = (Q \cup W) \cup R$ is a lens space.

This proves the first half of the lemma.

To finish we must identify the lens space coefficients of $P = P(p, 1)$. Let γ be the $1/2$ slope on the boundary of a solid torus V . Since γ crosses the meridian exactly twice, it bounds a Möbius band M in V . If we double V across its boundary, we obtain $S^2 \times S^1$. Note M doubles to give a Klein bottle. We now alter the double by opening it along an annulus neighbourhood of γ inside of ∂V and regluing with a twist to obtain the lens space L . Note the Klein bottle persists, so L is again a prism manifold. The image of the meridian of V under the p -fold twist about γ is the gluing slope, $(2p + 1)/4p$.

Thus $L(4p, 2p + 1)$ is a lens space, is a prism manifold, and has fundamental group of order $4p$. Thus $L(4p, 2p + 1) \cong P(p, 1)$, and we are done. ■

Lemma 2.10. *Suppose that q/p is a slope with s/r as a Farey neighbour. Then $P(p, q)$ is double covered by the lens space $L \cong L(2pq, ps + qr)$. The subgroup of $\pi_1(L)$ that is fixed by the deck group has order $2p$*

Proof. Recall that $T \times I$ is a double cover of $K \tilde{\times} I$. Let U and V be solid tori, whose disjoint union double covers the solid torus W . Thus the lens space

$$L = U \cup (T \times I) \cup V$$

double covers $P = W \cup (K \tilde{\times} I)$. Let α and β be simple closed curves in T , lifting a^2 and b in K . Thus $\langle \alpha, \beta \rangle$ gives a framing of T . The meridians of U and V have slopes q/p and $-q/p$ with respect to this framing. Fact 2.5 now implies that

$$L \cong L(2pq, ps + qr).$$

The above decomposition of L gives a presentation of $\pi_1(L)$. Abelianising, we obtain the following:

$$\pi_1(L) \cong \langle \alpha, \beta \mid p\alpha + q\beta = 0, p\alpha - q\beta = 0 \rangle$$

The elements correspond to the integer lattice points (up to translation) in the parallelogram with vertices at $(0, 0)$, $(p, -q)$, $(2p, 0)$, and (p, q) . The deck group fixes α

while sending β to β^{-1} . The fixed points under this action are the lattice points with second coordinate zero. So the fixed subgroup has order $2p$. ■

We use Lemma 2.10 to decide if a given manifold is a prism manifold $P(p, q)$, assuming that a double cover of the manifold is a lens space.

3. Normal and almost normal surfaces

Definition 3.1. Suppose that f is a two-simplex. An arc, properly embedded in f , is *normal* if it misses the vertices of f and has endpoints on distinct edges. ■

Definition 3.2. Suppose that Δ is a tetrahedron. A disc D , properly embedded in Δ and transverse to the edges of Δ , is

- a *triangle* if ∂D consists of three normal arcs;
- a *square* if ∂D consists of four normal arcs.

In either case D is a *normal disc*. See the left-hand sides of Figure 3.1. ■

Definition 3.3. Suppose that (M, \mathcal{T}) is a triangulated three-manifold. A surface S , properly embedded in M , is *normal* if, for each tetrahedron $\Delta \in \mathcal{T}$, the intersection $S \cap \Delta$ is a disjoint collection of triangles and squares. ■

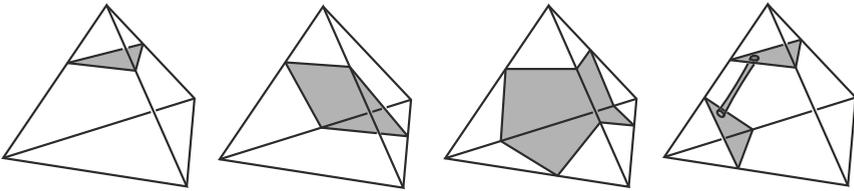


Figure 3.1. Left to right: triangle, square, octagon, tubed piece.

Definition 3.4. Suppose that Δ is a tetrahedron. A surface E , properly embedded in Δ and transverse to the edges of Δ , is

- an *octagon* if E is a disc with ∂E consisting of eight normal arcs;
- a *tubed piece* if E is an annulus obtained from two disjoint normal discs by attaching a tube that runs parallel to an arc of an edge of Δ .

In either case E is an *almost normal piece*. See the right-hand sides of Figure 3.1. ■

Definition 3.5. Suppose that (M, \mathcal{T}) is a triangulated three-manifold. A surface S , properly embedded in M , is *almost normal* if, for each tetrahedron $\Delta \in \mathcal{T}$, the intersection $S \cap \Delta$ is a disjoint collection of triangles and squares except for precisely

one tetrahedron, where the collection additionally contains exactly one almost normal piece. ■

Definition 3.6. A Heegaard surface S for a three-manifold M is *strongly irreducible* if it does not have disjoint compressing discs emanating from opposite sides of S . ■

We note that the genus one Heegaard surface for a lens space is strongly irreducible. We now have a result due to Stocking [28, Theorem 1], following work of Rubinstein.

Theorem 3.7. *Suppose that M is a closed, connected, oriented three-manifold equipped with a triangulation \mathcal{T} . Suppose that H is a strongly irreducible Heegaard surface for M . Then H is (ambiently) isotopic to a surface which is almost normal with respect to \mathcal{T} .* ■

Although almost normal surfaces are useful, they can sometimes be technically challenging to work with. We therefore apply the following result.

Proposition 3.8. *Suppose that S is almost normal with respect to a triangulation \mathcal{T} . Then S is isotopic to a surface which is normal with respect to the first barycentric subdivision $\mathcal{T}^{(1)}$.*

To prove this, we first require some lemmas.

Lemma 3.9. *Suppose that \mathcal{T} is a triangulation of a three-ball B . Suppose that for each tetrahedron Δ of \mathcal{T} the preimage of ∂B in Δ is either empty or is a single vertex, edge, or face. Suppose that F is a subsurface of ∂B so that ∂F is normal with respect to ∂B and intersects each edge of \mathcal{T} at most once. Then, after pushing the interior of F slightly into the interior of B , the resulting properly embedded surface F' is normal.*

Proof. Consider any tetrahedron Δ of \mathcal{T} . We divide into cases depending on the number of vertices of Δ contained in F .

If there are none, then $F' \cap \Delta$ is empty. Suppose that there is exactly one vertex v of Δ contained in F . Then $F' \cap \Delta$ is a normal triangle separating v from the remaining three vertices of Δ . This triangle meets ∂B in the set $(\partial F) \cap \Delta$. Suppose instead that there are exactly two vertices of Δ contained in F . Then there is an edge e of Δ contained in F . In this case $F' \cap \Delta$ is a normal square separating e from the remaining two vertices of Δ . Again, this square meets ∂B in the set $(\partial F) \cap \Delta$. Suppose instead that there are exactly three vertices of Δ contained in F . Then there is a face f of Δ contained in F . In this case $F' \cap \Delta$ is again a normal triangle, separating f from the final vertex of Δ . Also, this triangle is disjoint from ∂B . Note that, by assumption, we cannot have all four vertices of Δ contained in F . ■

Lemma 3.10. *Suppose that \mathcal{T} is a triangulation of a three-ball B . Suppose that for each tetrahedron Δ of \mathcal{T} the preimage of ∂B in Δ is either empty or is a single vertex, edge, or face. Let F_1, \dots, F_n be a collection of subsurfaces of ∂B . Suppose that each ∂F_i is normal and intersects each edge of \mathcal{T} at most once. Suppose also that for each $i \neq j$, ∂F_i and ∂F_j are disjoint and either $F_i \subset F_j$ or $F_j \subset F_i$. Then we may push the interiors of F_1, \dots, F_n slightly into the interior of B so that F' , the resulting union of surfaces, is properly embedded and normal.*

Proof. We apply the construction in the proof of Lemma 3.9 to each F_i , to form a normal surface F'_i . We note that when $i \neq j$, we can arrange for F'_i and F'_j to be disjoint, as follows. There is some F_i with the property that no other F_j lies inside it. For this F_i , form the resulting surface F'_i , which we can view as lying extremely close to ∂B . When we form the remaining surfaces F'_j , we can ensure that they are disjoint from this F'_i . In this way, the required surface is constructed recursively. ■

Proof of Proposition 3.8. We first specify the intersection between S and the edges of $\mathcal{T}^{(1)}$ lying in the one-skeleton of \mathcal{T} . By assumption, S has an almost normal piece P in some tetrahedron Δ of \mathcal{T} . The intersection between P and each one-simplex of Δ is at most two points. If it is exactly two points, then we arrange for these to lie in distinct edges of $\mathcal{T}^{(1)}$. We then arrange that the remaining points of intersection between S and the one-skeleton of \mathcal{T} are disjoint from the vertices of $\mathcal{T}^{(1)}$.

The intersection between S and each two-simplex F of \mathcal{T} consists of a collection of arcs that are normal with respect to \mathcal{T} . We may realise these arcs $F \cap S$ as a concatenation of normal arcs in the two-skeleton of $\mathcal{T}^{(1)}$, in such a way that each arc of $F \cap S$ intersects each edge of $\mathcal{T}^{(1)}$ at most once.

Thus, we have specified the intersection between S and the simplices of $\mathcal{T}^{(1)}$ lying in the two-skeleton of \mathcal{T} . The remainder of the three-manifold is a collection of tetrahedra of \mathcal{T} . Consider any such tetrahedron Δ . This inherits a triangulation from $\mathcal{T}^{(1)}$. We have already specified $S \cap \partial\Delta$. This is a collection of normal curves in $\partial\Delta$.

Suppose first that Δ does not contain the almost normal piece of S . Then $S \cap \Delta$ is a collection of triangles and squares in Δ . Their boundary is a collection C_1, \dots, C_n of normal curves in $\partial\Delta$. We need to specify, for each C_i , a subsurface F_i of $\partial\Delta$ that it bounds. For each C_i bounding a triangle in Δ , we pick F_i so that it contains a single vertex of Δ . The remaining C_i bound normal squares in Δ . We pick the F_i that these curves C_i bound so that they are nested. Then applying Lemma 3.10, we realise the discs bounded by these curves in Δ as normal surfaces with respect to $\mathcal{T}^{(1)}$.

Suppose now that Δ does contain the almost normal piece P of S . Then $S \cap \Delta$ is equal to P plus possibly some triangles and squares. In the case where P is an octagon, its boundary divides $\partial\Delta$ into two discs, and we pick one of these discs to

be the relevant F_i . In the case where P is obtained by tubing together two normal discs in Δ , the boundary curves of these discs cobound an annulus in $\partial\Delta$, and we set F_i to be this annulus. The remaining components of $S \cap \Delta$ are triangles and squares, and squares can only arise in the case where P is tubed. For each triangle with boundary C_j , we pick $F_j \subset \partial\Delta$ so that it contains a single vertex of Δ . For each square with boundary C_j , we pick $F_j \subset \partial\Delta$ so that it is disjoint from the annulus F_i considered above. Applying Lemma 3.10 again, we arrange for $S \cap \Delta$ to be normal with respect to $\mathcal{F}^{(1)}$. ■

4. Handle structures

Suppose that M is a compact, connected, oriented three-manifold. Suppose that \mathcal{H} is a handle decomposition of M . For example, we may obtain \mathcal{H} by taking the dual of a triangulation.

Definition 4.1. Suppose that S is a surface, properly embedded in M . We say that S is *standard* with respect to \mathcal{H} if the following properties hold:

- the intersection of S and any zero-handle $D^0 \times D^3$ is a disjoint union of properly embedded discs;
- the intersection of S and any one-handle $D^1 \times D^2$ is of the form $D^1 \times A$, where A is a disjoint union of arcs properly embedded in D^2 ;
- the intersection of S and any two-handle $D^2 \times D^1$ is of the form $D^2 \times P$, where P is a finite collection of points in the interior of D^1 ;
- the intersection of S and any three-handle $D^3 \times D^0$ is empty. ■

Definition 4.2. Let M be a closed three-manifold with a handle structure \mathcal{H} . A disc properly embedded in a zero-handle H_0 of \mathcal{H} is *normal* if

- its boundary lies within the union of the one-handles and two-handles;
- its intersection with each one-handle is a collection of arcs;
- it runs over each component of intersection between H_0 and the two-handles in at most one arc, and this arc respects the product structure on the two-handles.

A disc properly embedded in a one-handle H_1 of \mathcal{H} is *normal* if

- it respects the product structure on H_1 ;
- its boundary lies within the union of the zero-handles and two-handles;
- it runs over each component of intersection between H_1 and the two-handles in at most one arc. ■

Definition 4.3. Let M be a closed three-manifold with a handle structure \mathcal{H} . A surface properly embedded within M is *normal* if it is standard and its intersection with each zero-handle is normal. This implies that its intersection with each one-handle is also normal. ■

When S is a normal surface embedded in a closed triangulated three-manifold M , it naturally becomes a standard surface in the dual handle structure. When a three-manifold with a handle structure is decomposed along a standard surface S , the resulting three-manifold $M \setminus S$ inherits a handle structure. Thus, we deduce that when a closed triangulated three-manifold is cut along a normal surface, then $M \setminus S$ inherits a handle structure.

Our handle structures satisfy the following condition.

Definition 4.4. A handle structure of a three-manifold is *locally small* if the following conditions hold:

- the intersection between any zero-handle and the union of the one-handles consists of at most 4 discs; and
- the intersection between any one-handle and the union of the two-handles consists of at most 3 discs. ■

The handle structure that is dual to a triangulation is locally small. Moreover, when a three-manifold M with a locally small handle structure is cut along a normal surface S , the resulting handle structure is again locally small.

5. Core curves of solid tori

Over the next two sections, we prove Theorem 1.3. This is achieved using affine handle structures, first introduced in [14]. Here is the definition.

Definition 5.1. An *affine handle structure* on a three-manifold M is a handle structure where each zero-handle and one-handle is identified with a compact polyhedron in \mathbb{R}^3 , so that

- each face of each polyhedron is convex (but the polyhedron identified with a zero-handle need not be convex);
- whenever a zero-handle and one-handle intersect, each component of intersection is identified with a convex polygon in \mathbb{R}^2 , in such a way that the inclusion of this intersection into each handle is an affine map onto a face of the relevant polyhedron;
- for each zero-handle H_0 , each component of intersection with a two-handle, three-handle or ∂M is a union of faces of the polyhedron associated with H_0 ;

- the polyhedral structure on each one-handle is the product of a convex two-dimensional polygon and an interval. ■

Definition 5.2. Let \mathcal{T} be a triangulation of a compact three-manifold M and let \mathcal{H} be the dual handle structure. Then the *canonical affine structure* on \mathcal{H} realises each zero-handle as a truncated octahedron. Specifically, one realises each tetrahedron of \mathcal{T} as regular and euclidean with side length 1, and then one slices off the edges and the vertices to form a truncated octahedron. Each one-handle of \mathcal{H} corresponds to a face of \mathcal{T} that does not lie wholly in ∂M , and hence corresponds to a pair of hexagonal faces of the truncated octahedra that are identified. Thus, the one-handle is identified with the product of a hexagon and an interval. (See Figure 5.1.) ■

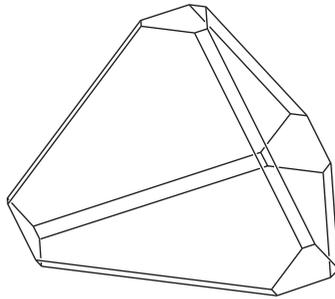


Figure 5.1. Each zero-handle is realised as a truncated octahedron. This is obtained from a tetrahedron by slicing off its vertices and edges.

The following theorem [14, Theorem 4.2] is the key result that goes into the proof of Theorem 1.3.

Theorem 5.3. *Let \mathcal{H} be a locally small, affine handle structure of the solid torus M . Then M has a core curve that intersects only the zero-handles and one-handles, that respects the product structure on the one-handles, that intersects each one-handle in at most 24 straight arcs, and that intersects each zero-handle in at most 48 arcs. Moreover, this collection of arcs in each zero-handle is parallel (as a collection) to a collection of arcs A in the boundary of the corresponding polyhedron. Finally, each component of A intersects each face of the polyhedron in at most 6 straight arcs. ■*

6. From affine handle structures to barycentric subdivisions

In this section, the goal is prove Theorem 1.3, which states that a core curve in a triangulated solid torus can be realised as a subcomplex in the triangulation's 51st

barycentric subdivision. Many of the technicalities of this section could have been avoided if we had aimed just for a subcomplex of the k th barycentric subdivision, for some universal but unspecified k independent of the triangulation.

Definition 6.1. A triangulation of a subset X of euclidean space is *straight* if the inclusion of each simplex into X is an affine map. ■

Definition 6.2. Two arcs properly embedded in a polygon and disjoint from its vertices are of the same *type* if there is an ambient isotopy taking one to the other, and which keeps the arcs disjoint from the vertices. ■

Note that a polygon with k sides can support at most $2k - 3$ disjoint straight arcs that are of distinct types.

Lemma 6.3. Let D be a euclidean polygon with a straight triangulation \mathcal{T} . Let α be a properly embedded straight arc in D . Then there is a realisation of the barycentric subdivision $\mathcal{T}^{(1)}$ as a straight triangulation of D that contains α as a subcomplex.

Proof. Since α is straight and \mathcal{T} is straight, the intersection between α and the interior of each one-simplex of \mathcal{T} is either all the interior of the one-simplex or at most one point. If α does intersect the interior of a one-simplex in a point, place a vertex of $\mathcal{T}^{(1)}$ at the point of intersection. Similarly, if α intersects the interior of a two-simplex, it does so in a single arc, and we place a vertex of $\mathcal{T}^{(1)}$ in the interior of this arc. Hence, α becomes simplicial in $\mathcal{T}^{(1)}$. (See Figure 6.1.) ■




Figure 6.1. Making an arc simplicial.

Induction then gives the following.

Lemma 6.4. Let D be a euclidean polygon with a straight triangulation \mathcal{T} . Let A be a union of k disjoint properly embedded straight arcs in D . Then there is a realisation of $\mathcal{T}^{(k)}$ as a straight triangulation of D that contains A as a subcomplex. ■

However, we can improve this in certain circumstances, as follows.

Lemma 6.5. Let D be a euclidean polygon with a straight triangulation \mathcal{T} . Let A be a union of at most 2^k disjoint properly embedded straight arcs in D , each with endpoints that are disjoint from the vertices of D , and that form at most two arc types.

Then there is a realisation of $\mathcal{T}^{(k+1)}$ as a straight triangulation of D that contains A as a subcomplex.

Proof. We prove this by induction on k . The case $k=0$ is the statement of Lemma 6.3. Let us prove the inductive step. Since the arcs fall into at most two types, there is one component α of A such that at most $|A|/2 \leq 2^{k-1}$ arcs lie on each side of it. By Lemma 6.3, there is a realisation of $\mathcal{T}^{(1)}$ as a straight triangulation of D that contains α as a subcomplex. Cutting D along α gives two euclidean polygons, each of which contains at most two arc types. Inductively, the intersection between A and these polygons may be made simplicial in $\mathcal{T}^{(k+1)}$. ■

Lemma 6.6. *Let D be a euclidean polygon with a straight triangulation \mathcal{T} . Let A be a union of at most 2^k disjoint properly embedded arcs in D , each with endpoints that are disjoint from the vertices of D , and that form at most $n \geq 2$ arc types. Then there is a realisation of $\mathcal{T}^{(n+k-1)}$ as a straight triangulation of D that contains A as a subcomplex.*

Proof. We prove this by induction on n . The induction starts with $n=2$, which is the content of Lemma 6.5. To prove the inductive step, suppose that $n > 2$. Pick an arc type of A and let α be a component of A of this type that is closest to other types of arcs. Make α simplicial in $\mathcal{T}^{(1)}$ using Lemma 6.3. Then cut along it, to give two euclidean polygons. In each, the arcs come in at most $n-1$ types. So, by induction, a further $n+k-2$ barycentric subdivisions suffice to make A simplicial. ■

Lemma 6.7. *Let D be a euclidean polygon with a straight triangulation \mathcal{T} . Let P be a set of at most 2^k points in D . Then there is realisation of $\mathcal{T}^{(k+1)}$ as a straight triangulation of D and that contains P in its vertex set.*

Proof. Pick a point x in D that is disjoint from P and that also does not lie on any line containing at least two point from P . We can pick a straight arc α through x , so that at most half the points of P lie on each side of α . This can be done as follows. Pick any straight arc α through x that misses P . If this has $|P|/2$ vertices on each side, then we have our desired arc. If not, then pick a transverse orientation on α that points to the side with more than $|P|/2$ points. Start to rotate α around x . By our general position hypothesis on x , at any given moment in time, the number of points on each side of α can jump by at most 1. By the time that α has rotated through angle π , the number of points on the side into which it points is less than $|P|/2$. So, at some stage, the arc contains a point of P and has fewer than $|P|/2$ points of P on either side of it. At this stage, we have our required arc α . By Lemma 6.3, we can make α simplicial in $\mathcal{T}^{(1)}$, and if $\alpha \cap P$ is non-empty, we can also make it a vertex. Cut along α and apply induction. ■

Lemma 6.8. *Let D be a euclidean polygon with a straight triangulation \mathcal{T} . Let A be a union of k disjoint straight arcs, with each endpoint being a vertex of \mathcal{T} or a point on ∂D , and with interior in the interior of D . Then there is realisation of $\mathcal{T}^{(k)}$ as a straight triangulation of D and that contains A as a subcomplex.*

Proof. We prove this by induction on k . Pick an arc α of A . Since α is straight, the intersection between α and the interior of each one-simplex of \mathcal{T} is either all the interior of the one-simplex or at most one point. If it does intersect the interior of this simplex in a point, place a vertex of $\mathcal{T}^{(1)}$ at the point of intersection. Similarly, if α intersects the interior of a two-simplex, it does so in a single arc, and we place a vertex of $\mathcal{T}^{(1)}$ in the interior of this arc. Hence, α becomes simplicial in $\mathcal{T}^{(1)}$. We then inductively deal with the remaining $k - 1$ arcs. ■

Lemma 6.9. *Let \mathcal{T} be a straight triangulation of a euclidean polyhedron P . Let C be a union of at most 2^k disjoint simple closed curves in ∂P that are simplicial in \mathcal{T} and that are topologically parallel in ∂P . Then there is realisation of $\mathcal{T}^{(2k)}$ as a straight triangulation of P such that C bounds a union of disjoint properly embedded discs in P that are simplicial in the triangulation.*

Proof. We prove this by induction on k . Since the components of C are all parallel in ∂P , there is some component C' of C that so that each component of $\partial P - C'$ contains at most half the components of C . We may realise a regular neighbourhood N of ∂P as a simplicial subset of $\mathcal{T}^{(2)}$. This is homeomorphic to $\partial P \times [0, 1]$, where $\partial P \times \{0\} = \partial P$. The annulus $C' \times [0, 1]$ may be realised as simplicial in $\mathcal{T}^{(2)}$. The curve $C' \times \{1\}$ bounds a disc in $\partial N - \partial P = \partial P \times \{1\}$. The union of this disc with the annulus $C' \times [0, 1]$ is one of the required discs. If we cut P along this disc, the result is two polyhedra P_1 and P_2 , with straight triangulations \mathcal{T}_1 and \mathcal{T}_2 . The intersection between $C - C'$ and each P_i consists of at most 2^{k-1} curves. By induction, these curves bound simplicial discs in $\mathcal{T}_i^{(2k-2)}$. Thus, C bounds simplicial discs in $\mathcal{T}^{(2k)}$. ■

Lemma 6.10. *Let \mathcal{T} be a straight triangulation of a euclidean polyhedron P . Let C be a union of at most 2^k disjoint simple closed curves in ∂P that are simplicial in \mathcal{T} . Suppose that the maximal number of pairwise non-parallel components of C is n . Then there is realisation of $\mathcal{T}^{(2k+2n-2)}$ as a straight triangulation of P such that C bounds a union of disjoint properly embedded discs in P that are simplicial in the triangulation.*

Proof. We prove this by induction on n . The induction starts with $n = 1$, which is the content of Lemma 6.9. To prove the inductive step, suppose that $n \geq 2$. Pick a curve type of C and let C' be a component of C of this type that is closest to other types of curves. As in the proof of Lemma 6.9, we may find a properly embedded disc

bounded by C' that is simplicial in $\mathcal{T}^{(2)}$. Cutting P along this disc gives two polyhedra, each of which inherits a triangulation. In each of these polyhedra, the maximal number of non-parallel components of $C - C'$ is at most $n - 1$. Thus, by induction, after barycentrically subdividing the triangulations of these polyhedra $(2k + 2n - 4)$ times, we obtain simplicial discs bounded by these curves. Hence, $\mathcal{T}^{(2k+2n-2)}$ contains simplicial discs bounded by C . ■

Lemma 6.11. *Let \mathcal{T} be a triangulation of a polyhedron P . Let A be a collection of disjoint simplicial arcs in ∂P . Let A' be the properly embedded arcs obtained by pushing the interior of A into the interior of P . Then, after an ambient isotopy supported in the interior of P , A' can be realised as simplicial in $\mathcal{T}^{(2)}$.*

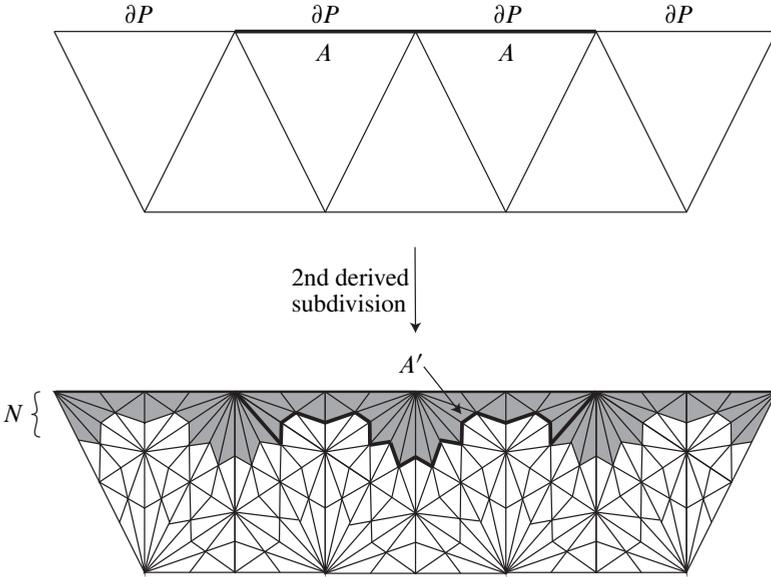


Figure 6.2. Pushing an arc into the interior.

Proof. Figure 6.2 gives a construction of A' . A regular neighbourhood N of ∂P is simplicial in $\mathcal{T}^{(2)}$. This is homeomorphic to $\partial P \times I$. Incident to the arcs A are simplicial discs of the form $A \times I$ in N . We then set $A' = \partial(A \times I) \setminus A$. ■

We now turn to the proof.

Theorem 1.3. *Let \mathcal{T} be a triangulation of the solid torus M . Then M contains a core curve that is a subcomplex of $\mathcal{T}^{(51)}$.*

Proof. We start with the triangulation \mathcal{T} of the solid torus M . Let \mathcal{H} be its dual handle structure. We give it its canonical affine structure, as in Definition 5.2. Each zero-handle is a truncated octahedron, which may be realised as a simplicial subset of the second derived subdivision of the tetrahedron of \mathcal{T} that contains it. Each one-handle of \mathcal{H} is realised as a product of a hexagon and an interval. We collapse this vertically onto its co-core, which is a hexagonal face of the two incident truncated octahedra.

We now apply Theorem 5.3, which provides a core curve C . The intersection between C and each hexagonal face is a collection of at most 24 points. These may be made simplicial after 6 barycentric subdivisions, by Lemma 6.7. Within each zero-handle, C is a union of at most 48 arcs, and these are simultaneously parallel to a collection of arcs A in the boundary of the truncated octahedron. We now make A simplicial. The intersection between A and each face of the truncated octahedron is at most $6 \times 48 = 288$ straight arcs. The ones that start and end on the boundary of the face come in at most 9 arc types and these can be made simplicial using at most 17 barycentric subdivisions, by Lemma 6.6. There are at most 24 arcs that have at least one endpoint not on the boundary of the face. We make these simplicial using at most 24 barycentric subdivisions using Lemma 6.8. Finally, we can push these arcs in the boundary of the truncated octahedron into the interior and make them simplicial, using at most 2 barycentric subdivisions, by Lemma 6.11. In total, we have used at most $2 + 6 + 17 + 24 + 2 = 51$ subdivisions. ■

7. Nicely embedded handle structures

Definition 7.1. Let \mathcal{T} be a triangulation of a compact three-manifold M , and let \mathcal{H} be the dual handle structure. A subset of a zero-handle or one-handle H of \mathcal{H} is *subnormal* if it is obtained from H by cutting along a collection of disjoint normal discs and then taking some of the resulting components. ■

The proof of Theorem 1.4 (in the case where the lens space M is not a prism manifold) proceeds by finding, within M , one of the solid tori V in its Heegaard splitting embedded in a nice way. More specifically, it has a handle structure where the union of the zero-handles and the one-handles is embedded in M in the following way.

Definition 7.2. Let M' be a handlebody embedded in a three-manifold M . Let A be a union of disjoint annuli in $\partial M'$. Let \mathcal{H}' and \mathcal{H} be handle structures for M' and M . We say that (\mathcal{H}', A) is *nicely embedded* in \mathcal{H} if the following hold:

- \mathcal{H}' has only zero-handles and one-handles;

- each zero-handle of \mathcal{H}' is a subnormal subset of a zero-handle of \mathcal{H} ;
- each one-handle of \mathcal{H}' is a subnormal subset of a one-handle of \mathcal{H} and has the same product structure;
- the intersection between the annuli A and any handle H' of \mathcal{H}' is a union of components of intersection between H' and handles of \mathcal{H} . ■

Definition 7.3. With notation as in the previous definition, we say that (\mathcal{H}', A) is (k, ℓ) -nicely embedded in \mathcal{H} , for non-negative integers k and ℓ , if it is nicely embedded and, in addition, the following hold:

- in any zero-handle of \mathcal{H} , at most k zero-handles of \mathcal{H}' lie between parallel normal discs in \mathcal{H} ;
- in any one-handle of \mathcal{H} , at most ℓ one-handles of \mathcal{H}' lie between parallel normal discs in \mathcal{H} . ■

Definition 7.4. Let M'' be a three-manifold embedded in another three-manifold M . Let \mathcal{H}'' and \mathcal{H} be handle structures for M'' and M . Let \mathcal{H}' be the handle structure just consisting of the zero-handles and one-handles of \mathcal{H}'' , and let A be the attaching annuli of the two-handles. We say that \mathcal{H}'' is (k, ℓ) -nicely embedded in \mathcal{H} , for non-negative integers k and ℓ , if (\mathcal{H}', A) is (k, ℓ) -nicely embedded in \mathcal{H} . ■

These definitions are designed to capture the essential properties of the handle structure that $M \setminus S$ inherits when S is a normal surface. More specifically, we have the following.

Lemma 7.5. *Suppose that M is a compact three-manifold with triangulation \mathcal{T} . Suppose that S is a normal surface, properly embedded in M . Let \mathcal{H} be the handle structure dual to \mathcal{T} and \mathcal{H}' be the handle structure that $M \setminus S$ inherits, but with the two-handles removed. Let A be their attaching annuli. Then (\mathcal{H}', A) is nicely embedded in \mathcal{H} . ■*

We typically collapse the one-handles of \mathcal{H}' vertically onto their co-cores. Thus, the underlying manifold of \mathcal{H}' becomes a collection of balls, which are just its zero-handles, glued along discs in their boundary.

Once we have such a (k, ℓ) -nice embedding, we get the following results.

Theorem 7.6. *Let M be a compact three-manifold with a triangulation \mathcal{T} . Let M' be a handlebody with a handle structure \mathcal{H}' , and let A be a union of disjoint annuli in $\partial M'$. Suppose that M' is embedded in M in such a way that (\mathcal{H}', A) is (k, ℓ) -nicely embedded in the dual of \mathcal{T} . Then we can arrange that the following are all simplicial subsets of $\mathcal{T}^{(m)}$:*

- each zero-handle of M' ;

- each one-handle of M' , vertically collapsed onto its co-core;
- the annuli A ;

where $m = 17 + 2\lceil \log_2(2k + 10) \rceil + \lceil \log_2(6 + 2\ell) \rceil + \lceil \log_2(4 + 2\ell) \rceil$.

Theorem 7.7. *Let M be a compact three-manifold with a triangulation \mathcal{T} . Let V be a solid torus with a handle structure \mathcal{H}' . Suppose that V is embedded in M in such a way that \mathcal{H}' is (k, ℓ) -nicely embedded in the dual of \mathcal{T} . Then there is a core curve of V that is a subcomplex of $\mathcal{T}^{(m+49)}$, where m is as in Theorem 7.6.*

Proof of Theorems 7.6 and 7.7. We start with the triangulation \mathcal{T} of M . Let \mathcal{H} be its dual handle structure. We give it its canonical affine structure, as in the proof of Theorem 1.3. Each zero-handle is a truncated octahedron, which may be realised as a simplicial subset of the second derived subdivision of the tetrahedron of \mathcal{T} that contains it. Each one-handle is the product of a hexagon and an interval, but it is collapsed onto its hexagonal co-core.

Our goal is to construct an affine handle structure on \mathcal{H}' . Thus, each zero-handle of \mathcal{H}' is given the structure of a euclidean polyhedron. We realise this as a polyhedron in the truncated octahedron of \mathcal{H}' that contains it and as a simplicial subset of a suitable iterated barycentric subdivision of \mathcal{T} .

Now, within each zero-handle H_0 of \mathcal{H} , every zero-handle of \mathcal{H}' is subnormal. These subnormal zero-handles are obtained from the truncated octahedron H_0 by cutting along normal triangles and squares. These are arranged into at most 4 triangle types and at most one square type. Our approach, in overview, is to arrange for the intersections between these normal discs and the one-handles of \mathcal{H} to be simplicial, then for the remainder of the boundary of these discs to be simplicial and then for the normal discs themselves to be simplicial.

Let us first focus on a one-handle of \mathcal{H} , which is a product of a hexagon X and an interval $[-1, 1]$. Within this one-handle, there are various one-handles of \mathcal{H}' . Only four possible one-handles of \mathcal{H}' in $X \times [-1, 1]$ do not lie between parallel normal discs. By assumption, at most ℓ one-handles do lie between parallel normal discs. These one-handles are therefore obtained from $X \times [-1, 1]$ by cutting along at most $6 + 2\ell$ normal discs and then possibly throwing away some components. The union of these discs is of the form $\beta \times [-1, 1]$ for normal arcs β in X . These arcs come in at most 3 types. Hence, by Lemma 6.6, after at most $2 + \lceil \log_2(6 + 2\ell) \rceil$ barycentric subdivisions, we make β simplicial in the triangulation of X .

Now consider each zero-handle H_0 of \mathcal{H} . We cut this handle along normal discs and then take some of the resulting components to get the subnormal zero-handles of \mathcal{H}' in H_0 . We are going to realise the boundary of these normal discs as simplicial in a suitable subdivision of the triangulation. We have already arranged for their intersection with the one-handles to be simplicial. Their intersection with each

two-handle consists of at most $4 + 2\ell$ arcs, all of the same arc type, and so by using $1 + \lceil \log_2(4 + 2\ell) \rceil$ further barycentric subdivisions, we can also arrange for these arcs to be simplicial, by Lemma 6.5.

The number of normal discs that we need to consider within H_0 is at most $2k + 10$. We now make these simplicial, using Lemma 6.10. At most $2\lceil \log_2(2k + 10) \rceil + 8$ barycentric subdivisions are required to do this.

Thus, each zero-handle of \mathcal{H}' is now a simplicial subset of $\mathcal{T}^{(m)}$, where m is as given in the statement of the theorem. Furthermore, when two zero-handles of \mathcal{H}' are joined by a one-handle, then we glue the simplicial subsets of $\mathcal{T}^{(m)}$ corresponding to these zero-handles along some faces. These can be arranged to be flat euclidean convex polygons with at most 6 sides. Thus, we have realised each one-handle, when vertically collapsed onto its co-core, as a simplicial subset of $\mathcal{T}^{(m)}$. Also, the components of intersection between the two-handles and the zero-handles and between the two-handles and the collapsed one-handles are simplicial. Thus, we have proved Theorem 7.6.

Let us now prove Theorem 7.7. So V is now a solid torus. Each zero-handle and one-handle of \mathcal{H}' has the structure of a euclidean polyhedron, as described above. This gives \mathcal{H}' an affine handle structure. We can therefore apply Theorem 5.3, which gives a core curve C of the solid torus V .

The intersection between C and any one-handle of \mathcal{H}' is at most 24 arcs, which respect the product structure on the handle. When the one-handle is vertically collapsed onto its co-core, these arcs become points. Using Lemma 6.7, we can make these vertices in the triangulation of the co-core of the one-handle after 6 barycentric subdivisions.

The intersection between C and each zero-handle of \mathcal{H}' is a trivial tangle. Moreover, we have control over the arcs in the boundary of the handle to which it is parallel. Specifically, these arcs are a union of straight arcs in each face of the handle. These arcs come in two types: those that start and end on the boundary of the face, and those that have at least one endpoint in C . In each face (that arises as a component of intersection with the one-handles), there are at most 24 of the latter type of arc. Hence, we need at most 24 barycentric subdivisions to make these simplicial, by Lemma 6.8. There are at most $48 \times 6 = 288$ arcs in each face that start and end on the boundary of the face and these come in at most 9 types. At most, 17 subdivisions are required to make these simplicial, by Lemma 6.6.

Now in each zero-handle, C runs parallel to these arcs, which have been made simplicial. Hence, using Lemma 6.11, two further subdivisions are required to make C simplicial.

The total number of barycentric subdivisions we have performed is at most $m + 49$. Thus, we have proved Theorem 7.7. ■

8. Parallellity bundles

In this section, we recall some material from [15, Section 5] about parallellity bundles for handle structures. Here we assume that M is a compact orientable three-manifold. We further assume that \mathcal{H} is a handle structure for M .

Definition 8.1. Suppose that γ is a simple closed curve, properly embedded in ∂M . We say that γ is *standard* with respect to \mathcal{H} if it satisfies the following properties:

- the curve γ is disjoint from the two-handles of \mathcal{H} ;
- for each one-handle $H = D^1 \times D^2$, there is a finite set $P \subset \partial D^2$ so that $\gamma \cap H = D^1 \times P$; and
- γ meets at least one one-handle. ■

Definition 8.2. Suppose that $S \subset \partial M$ is a subsurface. We say that \mathcal{H} is a *handle structure for the pair* (M, S) if

- \mathcal{H} is a handle structure for M ; and
- the boundary of S in ∂M is a union of standard curves for \mathcal{H} . ■

Definition 8.3. Suppose that \mathcal{H} is a handle structure for the pair (M, S) . Suppose that H is a zero-, one-, or two-handle of \mathcal{H} . We say that H is a *parallellity handle* if it admits a product structure $D^2 \times I$, so that

- $H \cap S = D^2 \times \partial I$; and
- for any other handle H' of \mathcal{H} every component of $H \cap H'$ is *vertical* in H : that is, of the form $\beta \times I$ where β is an arc in ∂D^2 . ■

Following the above definition, we typically regard parallellity handles as I -bundles over D^2 , their first coordinate.

Definition 8.4. The union of the parallellity handles in \mathcal{H} is the *parallellity bundle* for \mathcal{H} . ■

By [15, Lemma 5.3] the I -bundle structures on the parallellity handles agree where they intersect and so give an I -bundle structure on the parallellity bundle.

Definition 8.5. Suppose that \mathcal{B} is an I -bundle over a surface F . The resulting (∂I) -bundle over F is $\partial_h \mathcal{B}$, the *horizontal boundary* of \mathcal{B} . The resulting I -bundle over ∂F is $\partial_v \mathcal{B}$, the *vertical boundary* of \mathcal{B} . The components of the boundary of $\partial_h \mathcal{B}$ (which equals the boundary of $\partial_v \mathcal{B}$) are called the *corner curves* of \mathcal{B} . ■

Definition 8.6. Suppose that \mathcal{H} is a handle structure for the pair (M, S) . Suppose that \mathcal{B}^+ is a three-dimensional submanifold of M . We say that \mathcal{B}^+ is a *generalised parallellity bundle* if

- \mathcal{B}^+ is an I -bundle over a compact surface;
- the horizontal boundary of \mathcal{B}^+ is $\mathcal{B}^+ \cap S$;
- \mathcal{B}^+ is a union of handles of \mathcal{H} ;
- any handle in \mathcal{B}^+ that intersects the vertical boundary of \mathcal{B}^+ is a parallelity handle, where the I -bundle structure on the parallelity handle agrees with the I -bundle structure of \mathcal{B}^+ ;
- for any i -handle lying in \mathcal{B}^+ that is incident to j -handle, where $j > i$, the j -handle must also lie in \mathcal{B}^+ . ■

Note that the parallelity bundle \mathcal{B} is itself a generalised parallelity bundle.

Definition 8.7. We say that a generalised parallelity bundle \mathcal{B}^+ is *maximal* if \mathcal{B}^+ is not properly contained in another generalised parallelity bundle. ■

Note that there is always a maximal generalised parallelity bundle \mathcal{B}^+ that contains all of \mathcal{B} . However, the inclusion of \mathcal{B} into \mathcal{B}^+ need not respect the I -bundle structure on all components of \mathcal{B} .

Lemma 8.8. *Let \mathcal{B} be the parallelity bundle and let \mathcal{B}^+ be any maximal generalised parallelity bundle that contains \mathcal{B} . Then every corner curve of \mathcal{B}^+ is a corner curve of \mathcal{B} .*

Proof. By hypothesis, \mathcal{B}^+ contains \mathcal{B} . Hence, $\partial_h \mathcal{B}^+ = \mathcal{B}^+ \cap S$ contains $\partial_h \mathcal{B} = \mathcal{B} \cap S$. Suppose that γ is a corner curve of \mathcal{B}^+ . So γ is a component of ∂A for some component A of $\partial_v \mathcal{B}^+$. By definition, the handles of \mathcal{B}^+ incident to A are parallelity handles. The I -bundle structures on A and \mathcal{B} agree; also any parallelity handle incident to A lies in \mathcal{B} . Hence, A is a component of $\partial_v \mathcal{B}$. We deduce that γ is a corner curve of \mathcal{B} . ■

Definition 8.9. Suppose that G is an annulus, properly embedded in M , with boundary in S . Suppose that G' is an annulus in ∂M with $\partial G = \partial G'$. Suppose also that $G \cup G'$ bounds a three-manifold P such that

- either P is a parallelity region between G and G' or P lies in a three-ball;
- P is a non-empty union of handles;
- $\text{closure}(M - P)$ inherits a handle structure from \mathcal{H} ;
- any parallelity handle of \mathcal{H} that intersects P lies in P ;
- G is a vertical boundary component of a generalised parallelity bundle lying in P ;
- $G' \cap (\partial M - S)$ is either empty or a regular neighbourhood of a core curve of the annulus G' .

Removing the interiors of P and G' from M is called an *annular simplification*. ■

The resulting three-manifold M' , obtained from an annular simplification, is homeomorphic to the original manifold M . This holds even in the case where P is homeomorphic to the exterior of a non-trivial knot; in this case P lies in a three-ball in M . We now restate [15, Proposition 5.6].

Theorem 8.10. *Suppose that M is a compact orientable irreducible three-manifold. Suppose that F is an incompressible subsurface of ∂M . Let \mathcal{H} be a handle structure for (M, F) . Suppose that \mathcal{H} admits no annular simplification. Let \mathcal{B}^+ be any maximal generalised parallelity bundle in \mathcal{H} . Then the horizontal boundary of \mathcal{B}^+ is incompressible. ■*

9. Finding a simplicial core curve of a lens space

This section is devoted to the proof of Theorem 1.4: that is, for any triangulation \mathcal{T} of any lens space M (except $\mathbb{R}\mathbb{P}^3$) there is a relatively short curve with complement a solid torus or the orientation I -bundle over the Klein bottle.

Let \mathcal{H} be the handle structure of M that is dual to \mathcal{T} . Let S be an almost normal Heegaard torus in \mathcal{T} , which exists by Theorem 3.7. Cutting M along S gives two solid tori X_1 and X_2 . As explained in Section 4, these inherit handle structures \mathcal{H}_1 and \mathcal{H}_2 .

We now consider some particular situations where the proof of Theorem 1.4 is fairly straightforward. They highlight the approach that needs to be taken in the general case.

Suppose, as a special case, that one of \mathcal{H}_1 or \mathcal{H}_2 is $(0, 0)$ -nicely embedded within \mathcal{H} . We could then use Theorem 7.7 to find a core curve of one of the solid tori X_1 or X_2 that is simplicial in a suitable iterated barycentric subdivision of \mathcal{T} . But in general, the embeddings of \mathcal{H}_1 and \mathcal{H}_2 in \mathcal{H} are not $(0, 0)$ -nicely embedded, because when S contains two normal discs of the same type, the space between them becomes a zero-handle of \mathcal{H}_1 or \mathcal{H}_2 that violates the definition of a $(0, 0)$ -nice embedding.

For $i = 1$ and 2 , let \mathcal{B}_i be the parallelity bundle for \mathcal{H}_i . Suppose, as a different special case, that this is an I -bundle over a collection of discs, for $i = 1$ or 2 . Then we could create a new handle structure from \mathcal{H}_i by removing \mathcal{B}_i and replacing it by two-handles. This new handle structure is then $(0, 0)$ -nicely embedded in \mathcal{H} , and we could then apply Theorem 7.7.

Of course, though, there is no particular reason for \mathcal{B}_i to be I -bundles over discs. But according to Theorem 8.10, we can apply annular simplifications and then extend the parallelity bundle to a generalised parallelity bundle \mathcal{B}_i^+ with horizontal incompressible boundary. In fact, Theorem 8.10 is not immediately applicable, because if one sets F in that theorem to be all of the Heegaard surface S , then it is not incom-

pressible in the two solid tori X_1 and X_2 . But setting this aside for the moment, suppose that we could ensure that the generalised parallelity bundle \mathcal{B}_i^+ has horizontal incompressible boundary. It cannot be all of S , and so it is a union of disjoint discs and annuli. Hence, \mathcal{B}_i^+ consists of I -bundles over discs, annuli and Möbius bands. We can replace any I -bundles over discs by two-handles and remove any I -bundles over annuli using annular simplifications. Thus, if \mathcal{B}_i^+ contains no I -bundles over Möbius bands, then we end with a handle structure for one of the solid tori that is $(0, 0)$ -nicely embedded in \mathcal{H} . We could then apply Theorem 7.7.

To fix the problem that S is not incompressible in X_1 and X_2 , we work with the pair (X_i, F_i) , where F_i is a suitable subsurface of S . This is obtained from S by cutting along a curve or curves. One way of producing the required curves is via the following lemma.

Lemma 9.1. *Let V be a solid torus with a handle structure \mathcal{H} . Then there is a simple closed curve C in ∂V satisfying the following conditions:*

- *it is standard in \mathcal{H} ;*
- *it runs over each component of intersection between ∂V and the one-handles at most once;*
- *it is essential in ∂V and non-meridional.*

Suppose also that D is a union of disjoint discs in ∂V such that

- *it is a union of components of intersection between ∂V and the handles;*
- *if H and H' are handles of \mathcal{H} where H' has higher index than that of H , then whenever a component of $\partial V \cap H$ lies in D , so do all incident components of $\partial V \cap H'$.*

Then we may also ensure that C is disjoint from D .

Proof. Any essential curve in ∂V may be isotoped to a standard one, by first pushing it off the two-handles and then making it vertical in the one-handles. We consider all standard simple closed curves that are non-zero and non-meridional in $H_1(\partial V; \mathbb{Z}/2\mathbb{Z})$. We let C be such a curve that runs over the one-handles the fewest number of times. Then if it runs over a component of intersection between a one-handle and ∂V more than once, we can modify it to reduce this number by 2. This might create a disconnected one-manifold. But if so, then we just focus on one component. We can choose this component to be non-zero and non-meridional in $H_1(\partial V; \mathbb{Z}/2\mathbb{Z})$. Thus, under the assumption that C runs over the one-handles the fewest number of times, we deduce that C in fact runs over each component of intersection between the one-handles and ∂V at most once.

Consider now the case where there are also the discs D . Then we isotope any essential curve off D and make it standard. We consider such a curve that is non-zero

and non-meridional in $H_1(\partial V; \mathbb{Z}/2\mathbb{Z})$. Among all such curves disjoint from D , we let C be one that runs over the one-handles the fewest number of times. The same argument as above then applies. ■

Lemma 9.2. *Let M be a compact three-manifold with a handle structure \mathcal{H} . Let \mathcal{B} be the parallelity bundle for $(M, \partial M)$. Let C be a standard curve in ∂M that is disjoint from $\partial_h \mathcal{B}$. Let C' be three parallel copies of C , and let F be $\partial M \setminus N(C')$. Then the parallelity bundle for (M, F) is equal to \mathcal{B} .*

Proof. Every handle H of \mathcal{B} is, by definition, a parallelity handle for $(M, \partial M)$ and so satisfies $H \cap \partial M = H \cap \partial_h \mathcal{B}$. Since C is disjoint from $\partial_h \mathcal{B}$, it misses H , and therefore $H \cap F = H \cap \partial M$. So, H is a parallelity handle for (M, F) .

Now consider a parallelity handle H for (M, F) . If H is a two-handle, it is disjoint from the standard curves C' , and therefore $H \cap \partial M = H \cap F$. Hence, in this case, H is a parallelity handle for $(M, \partial M)$. Suppose now that H is a one-handle. It has the form $D^2 \times I$, where $H \cap F = D^2 \times \partial I$. Each component of intersection between H and any other handle has the form $\beta \times I$ for an arc β in ∂D^2 . There are two such components arising from the intersection between H and the incident zero-handles. There may be a further one or two components, arising from the intersection with the two-handles. If H does have two components of intersection with the two-handles, then $H \cap \partial M = H \cap F$, and therefore H is a parallelity handle for $(M, \partial M)$. On the other hand, if H has fewer than two components of intersection with the two-handles, then it intersects $\partial M \setminus F$ once or twice. Therefore, C' runs along the one-handle once or twice. But C' consists of three parallel copies of C , and therefore the number of times that it runs along this one-handle is a multiple of three. This is a contradiction. This completes the proof when H is a one-handle. Now suppose that H is a zero-handle. If it is disjoint from the one-handles, then C' misses it, since C' is standard. Hence, in this case, H is a parallelity handle for $(M, \partial M)$. On the other hand, if H intersects a one-handle, then this is also a parallelity handle for (M, F) and hence, as argued above, this one-handle has two components of intersection with the two-handles. Thus, $\partial D^2 \times I$ consists of intersections with one-handles and two-handles in an alternating fashion around the annulus $\partial D^2 \times I$. In particular, $H \cap \partial M = D^2 \times \partial I$, and therefore H is again a parallelity handle for $(M, \partial M)$. ■

We are now equipped to prove the theorem.

Theorem 1.4. *Let M be a lens space other than $\mathbb{R}\mathbb{P}^3$. Let \mathcal{T} be any triangulation of M . Then there is a simple closed curve C that is a subcomplex of $\mathcal{T}^{(86)}$, such that the exterior of C is either a solid torus or a twisted I -bundle over a Klein bottle.*

Proof. Let S be an almost normal Heegaard torus in \mathcal{T} , which exists by Theorem 3.7. By Proposition 3.8, S can be arranged to be normal in the barycentric subdivision $\mathcal{T}^{(1)}$.

Let \mathcal{H} be the handle structure of M that is dual to $\mathcal{T}^{(1)}$. Cutting M along S gives two solid tori X_1 and X_2 . These then inherit handle structures \mathcal{H}_1 and \mathcal{H}_2 .

Let \mathcal{B}_i be the parallelity bundle for $(X_i, \partial X_i)$ with handle structure \mathcal{H}_i . Then $\partial_v \mathcal{B}_i$ is a collection of properly embedded annuli in X_i . Their boundary curves are a (possibly empty) collection of essential curves on ∂X_i , each with the same slope α_i , together with some inessential curves on ∂X_i . We let $\alpha_i = \emptyset$ if there are no essential curves.

We consider three cases:

- (1) α_1 or α_2 is empty;
- (2) α_1 and α_2 are equal and non-empty;
- (3) α_1 and α_2 are distinct and non-empty.

Case (1). α_1 or α_2 is empty.

Say that α_i is empty. We claim that $\partial_h \mathcal{B}_i$ is a subsurface of S that lies in a collection of discs. Otherwise, $\partial_h \mathcal{B}_i$ contains a component F that is S minus some open discs. There cannot be another copy of F in S that is disjoint from F , and so F must be the horizontal boundary of a twisted I -bundle component of \mathcal{B}_i . The zero section of this I -bundle is therefore a non-orientable surface. By attaching annuli to its boundary that lie in $\partial_v \mathcal{B}_i$ and then capping off with discs, we obtain a closed non-orientable surface in the solid torus. This does not exist, which proves the claim.

We may take the discs D that contain $\partial_h \mathcal{B}_i$ to be as in Lemma 9.1. Specifically, each boundary component of $\partial_h \mathcal{B}_i$ bounds a disc in S , and we set D to be the union of these discs (which may be nested). Hence, there is a standard curve C_i in S that misses these discs, runs over each component of intersection between ∂X_i and the one-handles of \mathcal{H}_i at most once and is essential in ∂X_i and non-meridional. Let C'_i be three parallel copies of C_i and let $F_i = S \setminus N(C'_i)$. By Lemma 9.2, the parallelity bundle for (X_i, F_i) is exactly \mathcal{B}_i .

We now apply as many annular simplifications to \mathcal{H}_i as possible, forming a handle structure \mathcal{H}'_i for a solid torus isotopic to X'_i with subsurface F'_i of $\partial X'_i$. This process does not introduce parallelity handles. Thus, the horizontal boundary of the parallelity bundle \mathcal{B}'_i for \mathcal{H}'_i also lies in a union of disjoint discs in $\partial X'_i$. We now extend this parallelity bundle to a maximal generalised parallelity bundle \mathcal{B}_i^+ . By Lemma 8.8, the boundary curves of $\partial_h \mathcal{B}_i^+$ form a subset of the boundary curves of $\partial_h \mathcal{B}'_i$. They are therefore also inessential in $\partial X'_i$. By Theorem 8.10, $\partial_h \mathcal{B}_i^+$ is incompressible in X'_i . Hence, it is a union of discs. So, \mathcal{B}_i^+ consists of I -bundles over discs. Remove these discs and replace them by two-handles. The resulting handle structure is $(0, 0)$ -nicely embedded in \mathcal{H} . So, by Theorem 7.7, a core curve of X'_i is simplicial in $\mathcal{T}^{(79)}$. This proves the theorem in this case.

Case (2). α_1 and α_2 are equal and non-empty.

For some $i \in \{1, 2\}$, α_i is non-meridional in X_i , since M is not $S^2 \times S^1$. Fix some such i . Let C_i be a boundary component of $\partial_v \mathcal{B}_i$ with slope α_i . Then isotope C_i a little away from $\partial_h \mathcal{B}_i$ so that it becomes a standard curve in \mathcal{H}_i . Let C'_i be three parallel copies of C_i and let F_i be $S \setminus \setminus N(C'_i)$. We view \mathcal{H}_i as a handle structure for (X_i, F_i) . Again the parallelity bundle for (X_i, F_i) is \mathcal{B}_i , by Lemma 9.2. Perform as many annular simplification to \mathcal{H}_i as possible, forming the three-manifold X'_i with subsurface F'_i of $\partial X'_i$. Let \mathcal{B}'_i be its parallelity bundle. The boundary curves of its horizontal boundary therefore are inessential in $\partial X'_i$ or have slope α_i . Then, since α_i is non-meridional in X_i , Theorem 8.10 gives that we may extend \mathcal{B}'_i to a generalised parallelity bundle \mathcal{B}_i^+ with horizontal boundary that is incompressible in X'_i . The boundary curves of $\partial_h \mathcal{B}_i^+$ are inessential or have slope α_i , by Lemma 8.8.

Each component of \mathcal{B}_i^+ is an I -bundle over a disc, annulus or Möbius band. In fact, no component is an I -bundle over an annulus, for the following reason. A vertical boundary component of such an I -bundle is an incompressible annulus properly embedded in the solid torus X'_i . It is therefore boundary parallel in X'_i . By picking the vertical boundary component of the I -bundle appropriately, we may assume that the product region between it and an annulus in $\partial X'_i$ contains the I -bundle. Hence, we could have performed an annular simplification, contradicting our assumption that as many of these as possible were performed.

Suppose that no component of \mathcal{B}_i^+ is an I -bundle over a Möbius band. Then we can replace \mathcal{B}_i^+ by two-handles as in Case (1). The resulting handle structure is $(0, 0)$ -nicely embedded in \mathcal{H} . So, by Theorem 7.7, a core curve of X'_i is simplicial in $\mathcal{T}^{(79)}$.

So we are left with the case where some component of \mathcal{B}_i^+ is an I -bundle over a Möbius band. Then let $j = 3 - i$. Then α_j cannot be meridional in X_j , because we could patch a meridian disc for X_j onto a Möbius band in X_i to get an embedded projective plane. This would imply that M is $\mathbb{R}P^3$, contrary to assumption. Note that here we are using the fact that α_1 and α_2 are equal. Thus, the above argument gives a maximal generalised parallelity bundle \mathcal{B}_j^+ for (X'_j, F'_j) with $\partial_h \mathcal{B}_j^+$ incompressible in X'_j .

The only situation where the theorem is not proved in this case is when the generalised parallelity bundles \mathcal{B}_i^+ and \mathcal{B}_j^+ both contain components B_i and B_j that are I -bundles over Möbius bands. Then, α_i and α_j are the boundaries of Möbius bands in X_i and X_j and, after an isotopy, these can be glued to form an embedded Klein bottle in M . The exterior of this Klein bottle is a solid torus V . This is a general fact about Klein bottles in lens spaces (see the proof of Lemma 2.8). But it can be seen directly as follows. The exterior of B_i in X'_i is a solid torus which winds twice along X'_i . Moreover, $\partial_v B_i$ is longitudinal in this solid torus. Similarly, the exterior of B_j in X'_j

is a solid torus. If the two annuli $\partial X'_i \setminus \partial_h B_i$ and $\partial X'_j \setminus \partial_h B_j$ are isotoped to be equal, then when $X'_i \setminus B_i$ and $X'_j \setminus B_j$ are glued along this annulus, the result is a solid torus that is isotopic to V . Each boundary component of $\partial_v B_i$ is isotopic to a core curve of V . In this case, we take the required curve C to be one of these boundary components. Let A be $\partial_v \mathcal{B}_i^+$ lying in the boundary of the manifold $M' = X'_i \setminus \mathcal{B}_i^+$. Then (M', A) , with its inherited handle structure, is $(0, 0)$ -nicely embedded in \mathcal{H} . By Theorem 7.6, A is simplicial in $\mathcal{T}^{(30)}$. In particular, C , which is a boundary component of A is simplicial in $\mathcal{T}^{(30)}$, as required.

Case (3). α_1 and α_2 are distinct and non-empty.

For $i \in \{1, 2\}$, let C_i be some boundary curve of $\partial_v \mathcal{B}_i$ with slope α_i , isotoped a little so that it becomes a standard curve. Note that this isotopy pushes C_i into handles of \mathcal{H}_i that are not parallelity handles.

Now let F be $S \setminus N(C_1 \cup C_2)$. Since this lies in the complement of C_1 and C_2 , which are essential curves with the distinct slopes, F is a union of discs. In particular, it is incompressible in X_1 and X_2 .

Pick some i and perform as many annular simplifications to (X_i, F) as possible, giving a handle structure \mathcal{H}'_i for a pair (X'_i, F') that is isotopic to (X_i, F) . Let \mathcal{B}_i^+ be a maximal generalised parallelity bundle for (X'_i, F') that contains its parallelity bundle. By Theorem 8.10, the horizontal boundary of \mathcal{B}_i^+ is incompressible. Since it is a subsurface of F , it is a union of discs. Thus, \mathcal{B}_i^+ consists of I -bundles over discs. Replace each of these with a two-handle, giving a handle structure \mathcal{H}' .

We claim that \mathcal{H}' is $(10, 6)$ -nicely embedded in \mathcal{H} . Because \mathcal{H}' is obtained from \mathcal{H} by cutting along the normal surface S , it is (k, ℓ) -nicely embedded in \mathcal{H} for some k and ℓ , but we need to show why we can take $k = 10$ and $\ell = 6$. Consider a zero-handle H'_0 of \mathcal{H}' that lies between normally parallel discs. Let H_0 be the zero-handle of \mathcal{H} that contains it. Since H'_0 does not lie in \mathcal{B}_i^+ , it is not a parallelity handle for (X'_i, F') . Hence, it intersects C_1 or C_2 in S . Hence, in a component of $H_0 \setminus S$ incident to H'_0 , there must be a zero-handle of \mathcal{H}' that does not lie between parallel normal discs of S . Thus, at most 10 zero-handles of \mathcal{H}' in H_0 do not lie between parallel normal discs. This number 10 is twice the number of normal disc types of S that can simultaneously exist within H_0 . Therefore, we may set $k = 10$. A similar argument gives that $\ell = 6$ also suffices.

Since \mathcal{H}' is $(10, 6)$ -nicely embedded in \mathcal{H} , Theorem 7.7 gives that there is a core curve of X'_i that is simplicial in $\mathcal{T}^{(86)}$. This proves the theorem in this case. ■

We now use Theorem 1.3 and Theorem 1.4 to prove our main technical result.

Theorem 1.5. *Let M be a lens space, which is neither a prism manifold nor a copy of $\mathbb{R}P^3$. Let \mathcal{T} be any triangulation of M . Then the iterated barycentric sub-*

division $\mathcal{T}^{(86)}$ contains a core curve of M in its one-skeleton. Furthermore, $\mathcal{T}^{(139)}$ contains in its one-skeleton the union of the two core curves.

Proof. Let \mathcal{T} be a triangulation of a lens space M , other than a prism manifold or $\mathbb{R}P^3$. Note that the lens spaces that contain an embedded Klein bottle are exactly the prism manifolds by Lemma 2.8. So, by Theorem 1.4, there is a core curve C of M that is simplicial in $\mathcal{T}^{(86)}$. A regular neighbourhood of C is simplicial in $\mathcal{T}^{(88)}$. Removing the interior of this regular neighbourhood gives a triangulated solid torus. By Theorem 1.3, this contains a core curve that is simplicial in its 51st barycentric subdivision. Hence, M contains the union of its two core curves as a simplicial subset of $\mathcal{T}^{(139)}$. ■

10. Elliptic manifolds

A three-manifold M is *elliptic* if there is a subgroup $\Gamma \subset \mathrm{SO}(4)$, acting freely on the three-sphere, with $M \cong S^3/\Gamma$. Since the universal cover of M is the three-sphere, the fundamental group $\pi_1(M)$ is necessarily finite. Since all elements of $\mathrm{SO}(4)$ are orientation preserving, all elliptic manifolds are orientable. To give names to all of the elliptic manifolds we invoke a beautiful result of Seifert and Threlfall [29, p. 568]. See also [30, Corollary 4.4.11].

Theorem 10.1. *Every elliptic manifold admits a Seifert fibred structure.*

Recall that a Seifert fibred structure \mathcal{F} on a three-manifold M is a foliation by circles where each circle C has a neighbourhood U that is a fibred solid torus. The fibre C has *Seifert invariant* q/p if every fibre in $\mathcal{F}|_{\partial U}$ has slope q/p with respect to a framing $\langle \lambda, \mu \rangle$. Here μ is the meridian of U and p is necessarily non-zero. Note that replacing λ by $\lambda + \mu$ changes the slope of the fibre to be $(q - p)/p$.

We simplify the notation (M, \mathcal{F}) to just M when the foliation is understood. We say that a fibre $C \subset M$ is *generic* if it has integral Seifert invariant (in other words, $|p| = 1$) and *critical* otherwise. If we quotient all fibres of \mathcal{F} to points we get the *base orbifold* $B = M/S^1$, where critical fibres project to orbifold points.

The quotient induces a surjection $\pi_1(M) \rightarrow \pi_1^{\mathrm{orb}}(B)$ where the kernel is cyclic (possibly trivial) and generated by the generic fibre. Since $\pi_1(M) \cong \Gamma$ is finite, we deduce that $\pi_1^{\mathrm{orb}}(B)$ is also finite. Here, then, are the possibilities for B .

Lemma 10.2. *Suppose B is a closed two-dimensional orbifold (without mirrors). Then $\pi_1^{\mathrm{orb}}(B)$ is finite if and only if $\chi^{\mathrm{orb}}(B) > 0$. Thus B is one of the following:*

- *Cyclic:* $S^2, S^2(p), P^2, S^2(p, q)$;
- *Dihedral:* $P^2(r), S^2(2, 2, r)$, where $r \geq 2$;

- *Tetrahedral*: $S^2(2, 3, 3)$;
- *Octahedral*: $S^2(2, 3, 4)$;
- *Icosahedral*: $S^2(2, 3, 5)$.

The names indicate the orbifold fundamental group. ■

We call the last three the *platonic* orbifolds. In these cases we say the *type* of B is the orbifold fundamental group $\mathcal{P} = \pi_1^{\text{orb}}(B)$. Here \mathcal{P} is one of the three platonic groups A_4 , S_4 or A_5 , respectively.

Lemma 10.3. *Suppose (M, \mathcal{F}) is Seifert fibred and $B = M/S^1$ is the base orbifold. If B is cyclic or dihedral then M is $S^2 \times S^1$, a lens space or prism manifold.*

Proof. Suppose that $B = M/S^1$ is a two-sphere with at most two orbifold points. Let α be a loop in B cutting B into a pair of discs each with at most one marked point. Since M is orientable and since α is two-sided, the preimage of α in M is a torus T . The preimage of either component of $B - \text{int}(N(\alpha))$ is a fibred solid torus and thus M is a lens space or $S^2 \times S^1$.

Suppose that B is a projective plane with at most one orbifold point. Let α be a loop in B reversing orientation. So $B - \text{int}(N(\alpha))$ is a disc with at most one orbifold point. Since M is orientable, and since α is one-sided, the preimage of α in M is a Klein bottle. It follows that M is a prism manifold.

Suppose that $B = S^2(2, 2, r)$. Let α be an arc connecting the orbifold points of order 2. Again, the preimage of α is a Klein bottle and M is a prism manifold. ■

We say that a Seifert fibred space (M, \mathcal{F}) is *platonic* if $B = M/S^1$ is platonic. In this case, the *type* of M is the same as the type of B . Now Lemmas 10.2 and 10.3 allow us to coarsely name the elliptic manifolds.

Proposition 10.4. *Every elliptic manifold is either a lens space, a prism manifold, or a platonic manifold.* ■

To pin down an elliptic manifold precisely, we must discuss how the Seifert invariants of the critical fibres interact with the *Euler number*. Suppose, for this paragraph, that M is a platonic manifold. Suppose the Seifert invariants of the critical fibres C_i are q_i/p_i for $i = 1, 2, 3$. Following Orlik [20, p. 91], there is an integer q so that $\pi_1(M)$ has the following presentation:

$$\pi_1(M) = \langle a_1, a_2, a_3, b \mid [a_i, b] = 1, a_i^{p_i} = b^{-q_i}, a_1 a_2 a_3 = b^q \rangle.$$

Here $\langle b \rangle$ is the central subgroup and is generated by a generic fibre in M . Killing b gives $\pi_1^{\text{orb}}(B)$. The integer q is determined by the *Euler number* of (M, \mathcal{F}) :

$$e(M, \mathcal{F}) = -q - \sum \frac{q_i}{p_i}.$$

In a small abuse of notation we call the data $(q, q_1/p_1, q_2/p_2, q_3/p_3)$ the *Seifert invariants* of M . In the terminology of Orlik [20, p. 88], this manifold would be denoted by

$$\{q, (o_1, 0); (q_1, p_1), (q_2, p_2), (q_3, p_3)\}.$$

Suppose U_i is a fibred neighbourhood of the critical fibre $C_i \subset M$. If we change the framing of ∂U_i , say by replacing q_i/p_i with $(q_i + p_i)/p_i$, then in order to keep the Euler number the same we must also change the framing about a generic fibre, by replacing q with $q - 1$. We call this process *reframing*. Note also that there is an orientation-reversing homeomorphism between the manifolds having Seifert invariants

$$(q, q_1/p_1, \dots, q_n/p_n) \quad \text{and} \quad (-q, -q_1/p_1, \dots, -q_n/p_n).$$

Since we are not considering our manifolds as oriented in this paper, we view these two Seifert fibre spaces as equivalent.

We now record a very useful observation, which is essentially due to Seifert and Threlfall [29, p. 573].

Proposition 10.5. *Suppose M is a platonic manifold with*

$$B = M/S^1 = S^2(p_1, p_2, p_3).$$

Then the Seifert invariants of M can be recovered, up to reframing and reversing orientation, from B and the order of $H_1(M)$.

Proof. Abelianizing the fundamental group shows $H_1(M)$ has order

$$|qp_1p_2p_3 + q_1p_2p_3 + p_1q_2p_3 + p_1p_2q_3|.$$

A bit of modular arithmetic shows the data q, q_1, q_2, q_3 can be recovered, up to reframing and changing all their signs, from $B = M/S^1$ and the order of $H_1(M)$.

For example, suppose that B is $S^2(2, 3, 4)$. By applying an orientation-reversing homeomorphism if necessary, we may assume that $e(M, \mathcal{F}) \leq 0$. Its Seifert invariants are $(q, 1/2, q_2/3, q_3/4)$, where $q_2 \in \{1, 2\}$ and $q_3 \in \{1, 3\}$. Note that

$$|H_1(M)| = |24q + 12 + 8q_2 + 6q_3| = 24q + 12 + 8q_2 + 6q_3$$

since $e(M, \mathcal{F}) \leq 0$. We may compute q_2 using the fact that $|H_1(M)| \equiv 8q_2 \pmod{3}$, and we may compute q_3 using the fact that $|H_1(M)| \equiv 6q_3 + 12 \pmod{8}$. Then finally, q can be determined using the above equality for $|H_1(M)|$.

The cases of the other platonic manifolds are similar. ■

Remark 10.6. It is evident that the procedure given in the above proof runs in polynomial time as a function of the number of digits of $|H_1(M)|$.

Lemma 10.7. *Suppose M is a platonic manifold with base orbifold $B = M/S^1$. The map $\pi_1(M) \rightarrow \pi_1^{\text{orb}}(B)$ is the only surjection of the fundamental group of M onto a platonic group, up to post-composing by an automorphism of $\pi_1^{\text{orb}}(B)$.*

Proof. Suppose $G = \pi_1(M)$ and $\mathcal{P} = \pi_1^{\text{orb}}(B)$. Let $\rho: G \rightarrow \mathcal{P}$ be the associated surjection. Now suppose $\rho': G \rightarrow \mathcal{P}'$ is any surjection to a platonic group. Since \mathcal{P}' has trivial centre, the map ρ' kills the fibre of M . Thus ρ factors ρ' . However, there is no non-trivial surjective map between distinct platonic groups. Thus $\rho = \rho'$, perhaps after applying an automorphism of \mathcal{P} . ■

Proposition 10.8. *Suppose q/p and s/r are Farey neighbours, with $q \geq 2$. An orientable three-manifold M is homeomorphic to the prism manifold $P(p, q)$ if and only if M is double covered by the lens space $L = L(2pq, ps + qr)$ and the subgroup of $\pi_1(L)$ fixed by the deck group action has order $2p$.*

Proof. The forward direction is the statement of Lemma 2.10.

For the backwards direction let $G = \pi_1(L) \cong H_1(L)$. Since M is double covered by a lens space, the resolution of the spherical space form conjecture [21–23] implies that M is elliptic. Since $\pi_1(M)$ has an index two cyclic subgroup, namely G , we deduce that $\pi_1(M)$ cannot surject a platonic group.

Proposition 10.4 implies M is either a lens space or a prism manifold. However, in any double cover of a lens space, the deck group acts trivially on homology. We deduce that M is a prism manifold. According to Lemma 2.10, the coefficients p and q can be recovered from the order and index, respectively, of the subgroup of G that is fixed by the action of the deck group. ■

Proposition 10.9. *Suppose \mathcal{P} is one of the three platonic groups A_4, S_4 or A_5 . Let d be the order of \mathcal{P} . An orientable three-manifold M is homeomorphic to a platonic manifold of type \mathcal{P} if and only if it is d -fold covered by a lens space L with deck group isomorphic to \mathcal{P} . Moreover, the Seifert invariants of M can be recovered from the order of $H_1(M)$.*

Proof. For the forward direction, recall the fundamental group of M has the form

$$G = \pi_1(M) = \langle a_1, a_2, a_3, b \mid [a_i, b] = 1, a_i^{p_i} = b^{-q_i}, a_1 a_2 a_3 = b^q \rangle.$$

The subgroup $\langle b \rangle$ is central. The corresponding cover is an elliptic manifold with cyclic fundamental group, and so is a lens space. The deck group is isomorphic to $G/\langle b \rangle \cong \mathcal{P}$.

For the backwards direction, since M is covered by a lens space, the spherical space form conjecture (which is a consequence of the geometrisation conjecture, proved by Perelman [21–23]) implies that M is elliptic. Since the degree of the cover

equals the order of the deck group the covering is normal. Thus $\pi_1(M)$ is a cyclic extension of \mathcal{P} . We deduce M is not a lens space or prism manifold.

Proposition 10.4 implies M is a platonic manifold. Lemma 10.7 implies M has the type of \mathcal{P} . Finally, Proposition 10.5 states that the Seifert invariants can be recovered from the order of $H_1(M)$. ■

11. Certifying $T^2 \times I$ and $K^2 \tilde{\times} I$

As usual, we assume that any three-manifold M is given via a finite triangulation \mathcal{T} . The decision problem $\text{RECOGNISING } T^2 \times I$ takes \mathcal{T} as its input and it asks whether M is homeomorphic to $T^2 \times I$. The decision problem $\text{RECOGNISING } K^2 \tilde{\times} I$ is defined similarly. Both are dealt with by Haraway and Hoffman [8, Theorem 3.6].

Theorem 11.1. *The problems $\text{RECOGNISING } T^2 \times I$ and $\text{RECOGNISING } K^2 \tilde{\times} I$ are in NP.* ■

There is a subtle point to note here. Haraway and Hoffman, for their certificate for $\text{RECOGNISING } T^2 \times I$, rely on [16, Theorem 12.1]. There, given a handle structure \mathcal{H} of a sutured manifold (M, γ) , the theorem provides a certificate that (M, γ) is a product sutured manifold. In our setting, we would simply check that M had two toral boundary components T_0 and T_1 , say, and would assign the sutured manifold structure $R_-(M) = T_0$ and $R_+(M) = T_1$. Then [16, Theorem 12.1] provides a certificate that (M, \emptyset) is a product sutured manifold, which is equivalent to the statement that M is homeomorphic to $T^2 \times I$.

However, there is a gap in the published proof of [16, Theorem 12.1]. There one tacitly assumes that γ is non-empty. Nevertheless, there is a straightforward fix. We take as the certificate a non-separating annulus A in normal form with respect to the triangulation \mathcal{T} and with weight at most an exponential function of the number of tetrahedra of \mathcal{T} . Such an annulus exists by [16, Lemma 8.5, Theorem 8.3]. Then we form a handle structure on $M \setminus A$ where the number of handles is bounded above by a linear function of the number of tetrahedra of \mathcal{T} , using [16, Theorem 9.3]. This is a product sutured manifold if and only if M was a product; however now $M \setminus A$ has a non-empty collection of sutures and so the proof of [16, Theorem 12.1] applies.

We also note that there are alternative methods of certifying $T^2 \times I$. One is to use the following result.

Theorem 11.2. *Let M be a compact orientable three-manifold with two boundary components, both of which are tori. Let s_1, s_2, s_3 be slopes on one of the boundary components T of M that represent the three non-trivial elements of $H_1(T; \mathbb{Z}/2\mathbb{Z})$.*

Then M is homeomorphic to $T^2 \times I$ if and only if the three manifolds obtained by Dehn filling along s_1 , s_2 , and s_3 are all homeomorphic to a solid torus. ■

The proof is given inside the proof of [7, Theorem 11]. Thus we can certify that M is homeomorphic to $T^2 \times I$ as follows. We check that M has two boundary components, both of which are tori. Let T be one of these. We pick embedded normal one-manifolds C_1 , C_2 and C_3 that represent the three non-trivial elements of $H_1(T; \mathbb{Z}/2\mathbb{Z})$, with the property that each C_i intersects each triangle of T in at most one normal arc. Each C_i consists of some parallel essential simple closed curves in T , plus possibly some curves that bound discs in T . Let C'_i be one of these essential curves. The curves C'_1, C'_2, C'_3 again represent the three non-trivial elements of $H_1(T; \mathbb{Z}/2\mathbb{Z})$. We form triangulations for the three-manifolds obtained by Dehn filling along the slopes of C'_1, C'_2, C'_3 as follows. We barycentrically subdivide the triangulation of M once, so that each C'_i is simplicial. Then we attach (simplicially) a triangulated disc to M along C'_i . Finally, we attach a triangulated three-cell to complete the Dehn filling. Each of the resulting triangulations has at most 56 times as many tetrahedra as the original, given, triangulation of M . The final piece of our certificate is that these three triangulated three-manifolds are solid tori; for this we use [9, Corollary 2].

There is even a third method of providing a certificate for RECOGNISING $T^2 \times I$, going back to Schleimer's thesis [26, Chapter 6]. Let M be the given three-manifold equipped with the triangulation \mathcal{T} . We first check that M has the homology of T^2 . As in [25, Theorem 15.1], the first third of the certificate is a sequence $(\mathcal{T}_i, v(S_i))_{i=0}^n$, where

- $\mathcal{T}_0 = \mathcal{T}$;
- $v(S_i)$ is the normal vector of S_i , a fundamental non-vertex linking normal two-sphere in \mathcal{T}_i (for $i < n$);
- $v(S_n)$ is the zero-vector;
- \mathcal{T}_{i+1} is the result of *crushing* \mathcal{T}_i along S_i [25, Section 13]; and
- \mathcal{T}_n is *zero-efficient* [25, Definition 4.10].

We next certify three-sphere components of \mathcal{T}_n and discard them to obtain \mathcal{T}' . The final third of the certificate follows the plan of [26, Theorem 6.2.1]. We find a list $(T_i)_{i=0}^{2n}$ of disjoint tori in \mathcal{T}' . Each even torus T_{2k} (except for the first and last) is almost normal and *normalises via isotopy* to the normal tori $T_{2k\pm 1}$. The union $T_0 \cup T_{2n}$ is the frontier of the boundary of a small regular neighbourhood of ∂M , taken in \mathcal{T} . We require that T_0 and T_{2n} normalise via isotopy to T_1 and T_{2n-1} , respectively. That the T_i exist and have controlled weight is proved in [26, Chapter 6]. The normalisations are produced in polynomial time using the algorithm of [25, Theorem 12.1].

There are several possible certificates for $\text{RECOGNISING } K^2 \tilde{\times} I$. One is to exhibit a double cover \tilde{M} of the manifold M , to certify that \tilde{M} is $T^2 \times I$ and to check that M has a single torus boundary component. Specifically, the certificate is as follows:

- (1) a triangulation $\tilde{\mathcal{T}}$ of a three-manifold \tilde{M} ;
- (2) a simplicial involution ϕ of $\tilde{\mathcal{T}}$ with no fixed points;
- (3) a combinatorial isomorphism between $\tilde{\mathcal{T}}/\phi$ and \mathcal{T} ;
- (4) a certificate that \tilde{M} is homeomorphic to $T^2 \times I$.

We now check that ∂M is a single torus and verify the given certificate. This suffices because $K^2 \tilde{\times} I$ is the unique orientable three-manifold with a single toral boundary component and that is double covered by $T^2 \times I$ [6, Theorem 10.5].

12. Certificates for elliptic manifolds

In this section, we give our method for certifying whether a closed three-manifold is an elliptic manifold and, if it is, then which elliptic manifold it is. That is, we prove the following.

Theorem 1.1. *The problem ELLIPTIC MANIFOLD lies in NP.*

Theorem 1.2. *The problem NAMING ELLIPTIC lies in FNP.*

In what follows we assume that the three-manifold M is given via a finite triangulation \mathcal{T} .

12.1. Some problems in NP

In our proofs of Theorems 1.1 and 1.2, we rely on the fact that the following decision problems lie in NP:

- (1) $\text{RECOGNISING } S^3$: [25, Theorem 15.1] and [9, Theorem 2];
- (2) $\text{RECOGNISING } D^2 \times S^1$: [9, Corollary 2];
- (3) $\text{RECOGNISING } T^2 \times I$: [8, Theorem 3.6];
- (4) $\text{RECOGNISING } K^2 \tilde{\times} I$: [8, Theorem 3.6].

12.2. Some problems in P

In what follows we rely on the algorithm of Kannan and Bachem [10] which places a given integer matrix into Smith normal form. The running time and also the bit-size of the output are bounded above by polynomial functions of the bit-size of the input.

We also use the fact, due to Burton [4, Corollary 8], that it is decidable in polynomial time whether two triangulations of compact three-manifolds are combinatorially isomorphic. Burton proved this result by creating from a triangulation of a compact three-manifold an *isomorphism signature*, which gives a labelling of the simplices of the triangulation, and it has the property that two triangulations are combinatorially isomorphic if and only if their isomorphism signatures are equal [4, Theorem 7]. Furthermore, his algorithm also computes the automorphism group of a single triangulation in polynomial time.

We now provide various types of certificates for various types of elliptic manifolds. In each case, we give the certificate and explain how it is verified. The time to complete the verification is bounded above by a polynomial function of $|\mathcal{T}|$, the number of tetrahedra in \mathcal{T} . Finally, we prove that such a certificate exists if and only if M is the corresponding type of elliptic manifold.

In our certificates, we give more information than is strictly necessary. We do this in order to make the certificates easier to interpret by the reader of this paper. Some parts of the certificates could be removed, at the cost of requiring the checker of the certificate to do more work. We place an asterisk against those parts of the certificate that seem to crucial given the current state of knowledge.

12.3. Three-sphere

Recognising the three-sphere is discussed immediately above.

12.4. Real projective space

The certificate here is:

- (1) a triangulation $\tilde{\mathcal{T}}$ of a three-manifold \tilde{M} ;
- * (2) a certificate that \tilde{M} is the three-sphere;
- (3) a simplicial involution ϕ of $\tilde{\mathcal{T}}$ that has no fixed points; and
- (4) a simplicial isomorphism between $\tilde{\mathcal{T}}/\phi$ and \mathcal{T} .

The first, third, and fourth parts may be verified in polynomial time; we omit the details.

By the spherical space form conjecture [21–23] or [18], the manifold M has such a certificate if and only if it is $\mathbb{R}\mathbb{P}^3$. The output of the algorithm records whether the certificate has been verified, and in the function case, it also gives the Seifert data $(0, 1/2)$.

Only (2) appears to be necessary. Instead of (1), (3) and (4) being part of the certificate, one can proceed as follows. The verifier is given a triangulation \mathcal{T} of a three-manifold M . Using the algorithm of Kannan and Bachem, the verifier can

compute $H_1(M)$ in polynomial time and check that it is $\mathbb{Z}/2\mathbb{Z}$. The verifier can also construct the unique non-trivial homomorphism $\pi_1(M) \rightarrow \mathbb{Z}/2\mathbb{Z}$, and hence can build the unique double cover \tilde{M} and its lifted triangulation $\tilde{\mathcal{T}}$. Using Burton's isomorphism signature, the verifier can ensure that the simplices of this triangulation are labelled in a canonical fashion. The certificate in (2) applies to this labelled triangulation. Once the certificate in (2) is verified, the verifier can deduce that M is \mathbb{RP}^3 .

Remark 12.1. There is an interesting generalisation of the above. Suppose that $M = L(p, q)$ is given via a triangulation \mathcal{T} ; suppose that p is bounded by a uniform polynomial in $|\mathcal{T}|$. In this case, the verifier can construct the p -fold homology cover in polynomial time and also check that the deck group is cyclic. They then check a (provided) certificate that the cover is the three-sphere. This proves that M is a lens space. Finally, and more delicately, they verify the coefficient q in polynomial time using [12, Theorem 1.1].

12.5. Non-prism non- \mathbb{RP}^3 lens spaces

The certificate is as follows:

- * (1) a simplicial subset C of the one-skeleton of $\mathcal{T}^{(86)}$;
- (2) a triangulation \mathcal{T}' of a three-manifold X ; this will in fact be the exterior of C ;
- (3) a simplicial isomorphism between \mathcal{T}' and the triangulation that results from $\mathcal{T}^{(88)}$ by keeping only those simplices that are disjoint from C ;
- * (4) a certificate establishing that X is a solid torus;
- (5) simplicial one-chains λ and μ in $\mathcal{T}^{(88)}$; these will in fact be meridional and longitudinal curves on $\partial N(C)$;
- (6) positive integers p and q with $1 \leq q < p$; M will in fact be homeomorphic to $L(p, q)$.

To verify this certificate we must check the following:

- (1) that C is a circle;
- (2) that the mapping given in part (3) is a simplicial isomorphism;
- (3) that the certificate in part (4) shows that X is a solid torus.
- (4) that μ and λ are one-cycles that generate $H_1(\partial N(C))$;
- (5) that μ is trivial in $H_1(N(C))$ and that λ is a generator;
- (6) that the kernel of the homomorphism $H_1(\partial N(C)) \rightarrow H_1(X)$ (induced by inclusion) is generated by $p\lambda + q\mu$; and
- (7) that $(p, q) \neq (4p', 2p' \pm 1)$ for some integer p' .

Once verified this also, for the function case, records p and q . More specifically, since we are requiring that the output is the Seifert data for M , then this is $(0, q/p)$.

If we can verify the certificate, then M is the manifold $L(p, q)$. Conversely, suppose M is the lens space $L(p, q)$ and is neither a prism manifold nor $\mathbb{R}\mathbb{P}^3$. Then $(p, q) \neq (4p', 2p' \pm 1)$ by Lemma 2.9. By Theorem 1.5, the subdivision $\mathcal{T}^{(86)}$ contains a simplicial core curve C . Its regular neighbourhood $N(C)$ is simplicial in $\mathcal{T}^{(88)}$. An essential simple closed curve in $\partial N(C)$ that lies in the link of some vertex of C gives a simplicial meridian μ . An embedded simplicial arc in $\partial N(C)$ joining opposite sides of μ can then be closed up to form a simplicial longitude λ . Thus, the required certificate exists.

It can be verified in polynomial time, as a function of the number of tetrahedra in \mathcal{T} , for the following reasons. The subdivision $\mathcal{T}^{(88)}$ contains 24^{88} times as many tetrahedra as \mathcal{T} . Thus, steps (1)–(3) in the verification can be achieved in polynomial time. For (4)–(6), we note that the boundary maps in the chain groups for $N(C)$ and $\partial N(C)$ can be expressed as integer matrices, where the number of rows and columns is bounded by a constant multiple of the number of tetrahedra in $\mathcal{T}^{(88)}$ and where each column has L^1 norm at most 3. Thus, (4)–(6) can be verified in polynomial time using linear algebra, as in the Kannan-Bachem algorithm.

One can alternatively just use (1) and (4) in the certificate. Instead of (2) and (3), the verifier can compute the triangulation for the exterior X of C obtained from $\mathcal{T}^{(88)}$ by keeping only those simplices that are disjoint from C . Furthermore, by using Burton's isomorphism signature, the verifier can give the simplices of this triangulation a canonical labelling. Furthermore, the verifier can build the simplicial one-chains λ and μ and compute the integers p and q , using linear algebra. Note that, since they are computed in polynomial time, p and q necessarily have polynomial bit-size.

12.6. Prism lens spaces

Here, the certificate is either as in the case of non-prism non- $\mathbb{R}\mathbb{P}^3$ lens spaces (but where the integers p and q do satisfy $(p, q) = (4p', 2p' \pm 1)$ for some positive integer p') or the following:

- *(1) a simplicial subset C of the one-skeleton of $\mathcal{T}^{(86)}$;
- (2) a triangulation \mathcal{T}' of a three-manifold X ; this will in fact be the exterior of C ;
- *(3) a certificate that X is homeomorphic to $K^2 \tilde{\times} I$;
- (4) a triangulation $\tilde{\mathcal{T}}$ of a three-manifold \tilde{M} ; this will in fact be the double cover of M for which the inverse image of $K^2 \tilde{\times} I$ is a copy of $T^2 \times I$;
- (5) a simplicial involution ϕ of $\tilde{\mathcal{T}}$ that has no fixed points;
- (6) a simplicial isomorphism between \mathcal{T} and $\tilde{\mathcal{T}}/\langle\phi\rangle$;

- (7) a simplicial subset \tilde{C} of $\tilde{\mathcal{T}}^{(86)}$, partitioned into two subsets \tilde{C}_1 and \tilde{C}_2 ; this will in fact be the inverse image of C partitioned into its two components;
- (8) a triangulation $\tilde{\mathcal{T}}'$ of a three-manifold \tilde{X}' ; this will be the exterior of \tilde{C}_1 ;
- * (9) a certificate that \tilde{X}' is a solid torus;
- (10) simplicial one-chains λ and μ in $(\tilde{\mathcal{T}}')^{(88)}$; these will in fact be meridional and longitudinal curves on $\partial N(\tilde{C}_1)$; and
- (11) a positive integer p ; \tilde{M} will in fact be the lens space $L(2p, 1)$.

We check that C is a simple closed curve. We check that \mathcal{T}' is the triangulation obtained from $\mathcal{T}^{(88)}$ by removing those simplices incident to C . We verify the certificate that this is a copy of $K^2 \times I$, and so on. These imply that M is a prism manifold. Once we have checked that \tilde{M} is the lens space $L(2p, 1)$, we then check that ϕ_* acts trivially on $H_1(\tilde{M})$. By Lemma 2.10, this implies that M was the manifold $P(p, 1)$, which is indeed both a prism manifold and a lens space.

The existence of this certificate is a consequence of Theorem 1.4.

As in previous cases, some parts of this certificate can be dispensed with. One can alternatively compute $H_1(X)$ using the triangulation \mathcal{T}' provided in (2) and check that it is isomorphic to $\mathbb{Z} \times (\mathbb{Z}/2\mathbb{Z})$. One can then construct all double covers of M . One of these will be the cover \tilde{M} with triangulation $\tilde{\mathcal{T}}$. The inverse image of C in $\tilde{\mathcal{T}}^{(86)}$ is readily computed. There are two possible choices for the component C_1 . A triangulation of its exterior is readily computed, as are the one-chains λ and μ and the positive integer p .

Remark 12.2. Note that the above three cases, plus the case of S^3 , provide certificates for all lens spaces.

12.7. Prism non-lens spaces

The certificate:

- (1) a triangulation $\tilde{\mathcal{T}}$ of a three-manifold \tilde{M} ;
- (2) integers p, q, r, s where $p \geq 1$ and $q > 1$;
- * (3) a certificate that \tilde{M} is the lens space $L(2pq, ps + qr)$;
- (4) a simplicial involution ϕ of $\tilde{\mathcal{T}}$ that has no fixed points;
- (5) a simplicial isomorphism between \mathcal{T} and $\tilde{\mathcal{T}}/\langle\phi\rangle$.

To verify this certificate we must check the following:

- (1) that $ps - qr = 1$;
- (2) that the given certificate establishes that \tilde{M} is the lens space $L(2pq, ps + qr)$;
- (3) that the subgroup of $H_1(\tilde{M})$ that is fixed by ϕ has order $2p$;

- (4) that ϕ has no fixed points;
- (5) that the given map between \mathcal{T} and $\tilde{\mathcal{T}}/\langle\phi\rangle$ is a simplicial isomorphism.

Proposition 10.8 then implies that M is the prism manifold $P(p, q)$. This has two possible Seifert fibrations. The algorithm outputs the one with spherical base space, which has Seifert data $(0, 1/2, -1/2, q/p)$.

We need to show that if M is the manifold $P(p, q)$, then there is a certificate as above that can be verified in polynomial time. Proposition 10.8 gives that the prism manifold $P(p, q)$ has a double cover \tilde{M} that is the lens space $L(2pq, ps + qr)$. The triangulation \mathcal{T} lifts to a triangulation $\tilde{\mathcal{T}}$ for \tilde{M} . As explained in the previous subsections, there is a certificate that establishes that \tilde{M} is $L(2pq, ps + qr)$. As usual, $2pq$ has at most polynomially many digits (as a function of $|\tilde{\mathcal{T}}|$) as do p and q . Now, s/r can be any rational number so that $ps - qr = 1$. However, we may rechoose r , and s so that $0 \leq r \leq p$ and hence $|s| \leq |ps| = |1 + qr|$. Thus, it can be verified in polynomial time that $ps - qr = 1$. The remaining parts of the certificate may be checked in polynomial time. In particular, one uses linear algebra to verify that the subgroup of $H_1(\tilde{M})$ that is fixed by the covering involution ϕ has order $2p$.

Again certain parts of this certificate are not strictly necessary. Since the first homology of a prism manifold is generated by at most two elements, it has at most three double covers. For each such cover, one can compute its lifted triangulation and its isomorphism signature. One of these triangulations will be $\tilde{\mathcal{T}}$. If we proceed in this way, we do not need to be provided with (4) and (5) in the certificate. The order of the subgroup of $H_1(\tilde{M})$, fixed by the covering transformation, can be computed in polynomial time; this gives p . Since we have established that lens space recognition is in FNP, when this is applied to \tilde{M} , the integers $2pq$ and $ps + qr$ are given as the lens space coefficients; hence these do not need to be provided as part of a certificate. Once we have p and $2pq$, we can compute q . As observed above, s and r can be any integers satisfying $ps - qr = 1$ and hence can be computed using the euclidean algorithm, and can be arranged so that $0 \leq r \leq p$ and $|s| \leq |1 + qr|$. Then $ps + qr$ can be computed and we can check whether it equals the second integer that is output from the recognition of \tilde{M} as a lens space.

12.8. Platonic manifolds

The certificate:

- *(1) a triangulation $\tilde{\mathcal{T}}$ of a three-manifold \tilde{M} ;
- *(2) a group \mathcal{P} of simplicial isomorphisms of $\tilde{\mathcal{T}}$ that acts freely;
- (3) a simplicial isomorphism between $\tilde{\mathcal{T}}/\mathcal{P}$ and \mathcal{T} ;
- (4) an isomorphism between \mathcal{P} and one of the platonic groups A_4 , S_4 or A_5 ;

- * (5) a certificate that certifies that \tilde{M} is a lens space using one of the certificates described above in Sections 12.4, 12.5 or 12.6.

To verify this certificate we must check the following:

- (1) that the given isomorphism between \mathcal{P} and one of the platonic groups is indeed an isomorphism;
- (2) that \mathcal{P} acts freely on $\tilde{\mathcal{T}}$;
- (3) that the given simplicial isomorphism between $\tilde{\mathcal{T}}/\mathcal{P}$ and \mathcal{T} is indeed a simplicial isomorphism;
- (4) that the given certificate establishes that \tilde{M} is homeomorphic to a lens space;
- (5) the resulting Seifert invariants of M , using Proposition 10.5.

It is clear that if there is such a certificate, then M is regularly covered by a lens space, with deck group given by the platonic group \mathcal{P} . Hence, by Proposition 10.9, M is a platonic manifold with type \mathcal{P} . Conversely, if M is such a manifold, then by Proposition 10.9, there is such a finite cover, and hence the certificate exists. It can be verified in polynomial time.

It is clear that (3) and (4) are not actually needed as part of the certificate. The simplicial isomorphism in (3) can be computed using Burton's algorithm and it can readily be checked that the given group \mathcal{P} of covering transformations is isomorphic to one of A_4 , S_4 or A_5 .

An alternative certificate, instead of giving (1) and (2), could give a surjective homomorphism from $\pi_1(M)$ to \mathcal{P} , the desired platonic group. The verifier could then build the cover $\tilde{\mathcal{T}}$ and its deck group directly, in polynomial time.

Acknowledgements. We thank the referee for their detailed and insightful comments. We also thank the organisers of the conference on Geometric Topology of Knots at the University of Pisa, where this work was initiated.

Funding. M. L. was partially supported by the Engineering and Physical Sciences Research Council (grant number EP/Y004256/1).

References

- [1] M. Aschenbrenner, S. Friedl, and H. Wilton, [Decision problems for 3-manifolds and their fundamental groups](#). In *Interactions between low-dimensional topology and mapping class groups*, pp. 201–236, Geom. Topol. Monogr. 19, Geometry and Topology, Coventry, 2015 Zbl 1336.57001 MR 3609909
- [2] F. Bonahon and J.-P. Otal, [Scindements de Heegaard des espaces lenticulaires](#). *Ann. Sci. École Norm. Sup. (4)* **16** (1983), no. 3, 451–466 Zbl 0545.57002 MR 0740078

- [3] G. E. Bredon and J. W. Wood, [Non-orientable surfaces in orientable 3-manifolds](#). *Invent. Math.* **7** (1969), 83–110 Zbl [0175.20504](#) MR [0246312](#)
- [4] B. A. Burton, [The Pachner graph and the simplification of 3-sphere triangulations](#). In *Computational Geometry (SCG'11)*, pp. 153–162, Association for Computing Machinery, New York, NY, 2011 Zbl [1283.05065](#) MR [2919606](#)
- [5] H. Geiges and N. Thies, [Klein bottles in lens spaces](#). *Involve* **16** (2023), no. 4, 621–636 Zbl [1548.57052](#) MR [4662620](#)
- [6] J. Hempel, [3-manifolds](#). Reprint of the 1976 original, AMS Chelsea Publishing, Providence, RI, 2004 Zbl [1058.57001](#) MR [2098385](#)
- [7] R. C. Haraway III, [Determining hyperbolicity of compact orientable 3-manifolds with torus boundary](#). *J. Comput. Geom.* **11** (2020), no. 1, 125–136 Zbl [1485.57017](#) MR [4117666](#)
- [8] R. C. Haraway III and N. R. Hoffman, [On the complexity of cusped non-hyperbolicity](#). [v1] 2019, [v3] 2022, arXiv:[1907.01675v3](#)
- [9] S. V. Ivanov, [The computational complexity of basic decision problems in 3-dimensional topology](#). *Geom. Dedicata* **131** (2008), 1–26 Zbl [0175.20504](#) MR [2369189](#)
- [10] R. Kannan and A. Bachem, [Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix](#). *SIAM J. Comput.* **8** (1979), no. 4, 499–507 Zbl [0446.65015](#) MR [0573842](#)
- [11] H. Kneser, [Geschlossene Flächen in dreidimensionalen Mannigfaltigkeiten](#). *Naturwiss.* **16** (1928), 973 Zbl [54.0644.05](#)
- [12] G. Kuperberg, [Identifying lens spaces in polynomial time](#). *Algebr. Geom. Topol.* **18** (2018), no. 2, 767–778 Zbl [1432.57034](#) MR [3773737](#)
- [13] G. Kuperberg, [Algorithmic homeomorphism of 3-manifolds as a corollary of geometrization](#). *Pacific J. Math.* **301** (2019), no. 1, 189–241 Zbl [1434.57013](#) MR [4007377](#)
- [14] M. Lackenby, [Core curves of triangulated solid tori](#). *Trans. Amer. Math. Soc.* **366** (2014), no. 11, 6027–6050 Zbl [1321.57024](#) MR [3256192](#)
- [15] M. Lackenby, [The crossing number of composite knots](#). *J. Topol.* **2** (2009), no. 4, 747–768 Zbl [1190.57003](#) MR [2574742](#)
- [16] M. Lackenby, [The efficient certification of knottedness and Thurston norm](#). *Adv. Math.* **387** (2021), article no. 107796 Zbl [1520.57007](#) MR [4274879](#)
- [17] M. Lackenby and S. Schleimer, [Lens space recognition is in NP](#). In *Triangulations, Oberwolfach Rep.* **9** (2012), no. 2, 1421–1424
- [18] G. R. Livesay, [Fixed point free involutions on the 3-sphere](#). *Ann. of Math. (2)* **72** (1960), 603–611 Zbl [0096.17302](#) MR [0116343](#)
- [19] R. Myers, [Free involutions on lens spaces](#). *Topology* **20** (1981), no. 3, 313–318 Zbl [0508.57032](#) MR [0608603](#)
- [20] P. Orlik, [Seifert manifolds](#). Lect. Notes Math. 291, Springer, Berlin-New York, 1972 Zbl [0263.57001](#) MR [0426001](#)
- [21] G. Perelman, [The entropy formula for the Ricci flow and its geometric applications](#). 2002, arXiv:[math/0211159v1](#)
- [22] G. Perelman, [Finite extinction time for the solutions to the Ricci flow on certain three-manifolds](#). 2003, arXiv:[math/0307245v1](#)

- [23] G. Perelman, Ricci flow with surgery on three-manifolds. 2003, arXiv:[math/0303109v1](https://arxiv.org/abs/math/0303109v1)
- [24] K. Reidemeister, [Homotopieringe und Linsenräume](#). *Abh. Math. Sem. Univ. Hamburg* **11** (1935), no. 1, 102–109 Zbl [0011.32404](#) MR [3069647](#)
- [25] S. Schleimer, [Sphere recognition lies in NP](#). In *Low-dimensional and symplectic topology*, pp. 183–213, Proc. Sympos. Pure Math. 82, American Mathematical Society, Providence, RI, 2011 Zbl [1250.57024](#) MR [2768660](#)
- [26] S. Schleimer, Almost normal Heegaard splittings. PhD thesis, University of California, Berkeley, 2001
- [27] P. Scott and H. Short, [The homeomorphism problem for closed 3-manifolds](#). *Algebr. Geom. Topol.* **14** (2014), no. 4, 2431–2444 Zbl [1311.57025](#) MR [3331689](#)
- [28] M. Stocking, [Almost normal surfaces in 3-manifolds](#). *Trans. Amer. Math. Soc.* **352** (2000), no. 1, 171–207 Zbl [0933.57016](#) MR [1491877](#)
- [29] W. Threlfall and H. Seifert, [Topologische Untersuchung der Diskontinuitätsbereiche endlicher Bewegungsgruppen des dreidimensionalen sphärischen Raumes \(Schluß\)](#). *Math. Ann.* **107** (1933), no. 1, 543–586 Zbl [58.1203.01](#) MR [1512817](#)
- [30] W. P. Thurston, *Three-dimensional geometry and topology. Vol. 1*. Princeton Math. Ser. 35, Princeton University Press, Princeton, NJ, 1997 Zbl [0873.57001](#) MR [1435975](#)

Received 21 October 2022.

Marc Lackenby

Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK;
marc.lackenby@maths.ox.ac.uk

Saul Schleimer

Mathematics Institute, University of Warwick, Coventry CV4 7AL, UK;
s.schleimer@warwick.ac.uk