Algebraic & Geometric Topology 18 (2018) 767–778

Identifying lens spaces in polynomial time

GREG KUPERBERG

We show that if a closed, oriented 3–manifold M is promised to be homeomorphic to a lens space $L(n,k)$ with n and k unknown, then we can compute both n and k in polynomial time in the size of the triangulation of M. The tricky part is the parameter k. The idea of the algorithm is to calculate Reidemeister torsion using numerical analysis over the complex numbers, rather than working directly in a cyclotomic field.

57M27; 65G30, 68Q15, 68W01

1 Introduction

The algorithmic problem of distinguishing or classifying closed d–dimensional manifolds is easy when $d \leq 2$, provably impossible when $d \geq 4$ (see Section 7.1 of Poonen [16]), and recursive when $d = 3$ (see Kuperberg [8]). The remaining question is how efficiently we can distinguish closed 3–manifolds; or whether we can distinguish them efficiently with one or another form of help. One small but interesting part of this question is the case of lens spaces. If M is a closed, oriented 3–manifold, conventionally given by a triangulation, then is it a lens space? If so, which one? In this article, we show that at least the second question has an efficient algorithm.

Theorem 1.1 *Suppose that M is a closed, oriented 3–manifold given by a triangulation with t tetrahedra, and that we are promised that $M \cong L(n,k)$ is a lens space with n and k unknown. Then n and k can be computed in deterministic polynomial time in t.*

The motivation for our result is a recent result announced by Lackenby and Schleimer [13] to both recognize whether M is a lens space, and if so which one, in the complexity class FNP. In other words, they provide a deterministic algorithm (a verifier) with the help of a prover who asserts the answer and provides a certificate that it is correct. Thus, Theorem 1.1 implies that in the Lackenby–Schleimer result, it is enough for the prover to only provide a certificate that M is a lens space at all, which is simpler.

Published: 12 March 2018 DOI: 10.2140/agt.2018.18.767

According to Lackenby and Schleimer, the certificate can be a Heegaard torus which is almost normal relative to the triangulation of M.

Recall that the standard lens space $L(n,k)$ is constructed by gluing the top hemisphere of a ball, often imagined as a convex dihedron or "lens", to the bottom hemisphere with a rotation of $2\pi k/n$. The calculation of n is reasonably standard, because if $M \cong L(n,k)$, then we can calculate the homology $H_1(M) \cong \mathbb{Z}/n$ in polynomial time using a version of the Smith normal form algorithm; see Kannan and Bachem [6]. The second parameter k is more subtle. We can take it to be a prime residue $k \in (\mathbb{Z}/n)^\times$. Reidemeister [17] showed that

$$L(n,k_1) \cong L(n,k_2)$$

as oriented 3–manifolds if and only if $k_1 = k_2$ or $k_1 = 1/k_2$.

In another respect, both parameters are more subtle than one might expect. Suppose that $M \cong L(n,k)$ has t tetrahedra. In the most standard (generalized) triangulation of $L(n,k)$, we have $n = t$. But there are other families of triangulated manifolds $M \cong L(n,k)$ such that n is exponential in t, and with exponentially many values of k for specific values of n. See Section 2. If we can be promised a polynomial bound on n itself rather than merely a polynomial bound on its digits, then it is easier to calculate k, because we can directly follow Reidemeister's method by computing the Reidemeister torsion Δ of M (endowed with a suitable local system of coefficients) in the cyclotomic ring $\mathbb{Z}[\zeta_n]$ or its fraction field $\mathbb{Q}(\zeta_n)$, where ζ_n is a primitive n^{th} root of unity.

The idea of our proof of Theorem 1.1 is to approximately compute the Reidemeister torsion using numerical analysis over the complex numbers $\mathbb{C}$. If we let $\zeta_n = \exp(2\pi i/n)$, the result is a sparse polynomial expression

$$\Delta = \zeta_n^c(1-\zeta_n^a)(1-\zeta_n^b) \in \mathbb{C}.$$

In order to establish a polynomial-time algorithm, we want a polynomial upper bound on the digits of precision of an approximation to Δ that we need to resolve the exponents a, b and c. We also need an algorithm to calculate those exponents. More precisely, the precision bound needs to be polynomial in t and thus polynomial in $\log n$. According to the answers to a MathOverflow question posed by Terry Tao [21], even the first part is unknown for general sparse sums of roots of unity. A bound is known for sums with at most four terms; see Myerson [14]. The unproven behavior of sparse sums of powers of ζ_n can be circumvented by making ζ_n a randomly chosen primitive n^{th} root of unity rather than specifically $\exp(2\pi i/n)$, but this would still leave the difficulty of finding an efficient algorithm to calculate the exponents. In the special case of the quantity Δ that we need, the precision problem has a good bound for any choice of ζ_n, and there is an efficient algorithm to calculate the exponents, based in part on an answer to the author's question in MathOverflow [11]. We include this argument in full in this paper. (The answer was posted under a pseudonym and is still anonymous as of the end of 2017.)

In a previous version of this paper [10], the author found a weaker version of Theorem 1.1 with a quantum polynomial-time algorithm, ie an algorithm in BQP; see Nielsen and Chuang [15]. The idea then was to replace $\mathbb{Z}[\zeta_n]$ with a quotient field $\mathbb{Z}/p$, where p is a prime which is congruent to 1 mod n. Then the Reidemeister torsion calculation reduces to the discrete logarithm problem, which can be solved with Shor's algorithm [20]. A quantum algorithm which is faster than any competing classical algorithm is always interesting, but in this case the author later noticed that there is a fast classical algorithm after all.

The question remains whether there is a competitive quantum algorithm for any natural question in 3–manifold topology. This is a natural thing to look for, since for instance it is known that unknottedness is the complexity class $\mathsf{NP} \cap \mathsf{coNP}$; see Hass, Lagarias and Pippenger [5], Kuperberg [7] and Lackenby [12]. (See the Complexity Zoo [1] for a survey of computational complexity classes.) While $\mathsf{NP} \cap \mathsf{coNP}$ is thought to neither contain nor be contained in quantum polynomial time BQP, some key problems (such as discrete logarithm) are known to be in both of them. Aharonov, Jones and Landau [2] give an algorithm to approximate the Jones polynomial of a knot at a principal root of unity; this algorithm also has a version for 3–manifolds; see Garnerone, Marzuoli and Rasetti [4]. However, the approximation is exponentially poor: any fair approximation that could be useful for geometric topology is #P–hard; see Kuperberg [9].

Acknowledgments

The author would like to thank an anonymous MathOverflow user for help with part of the calculation. The author was partly supported by NSF grants CCF-1319245 and CCF-1716990.

2 Large lens spaces with small triangulations

In this section, we will construct lens spaces $M \cong L(n,k)$ where n is much larger than the number of tetrahedra t, and k has many possible values. The manifolds

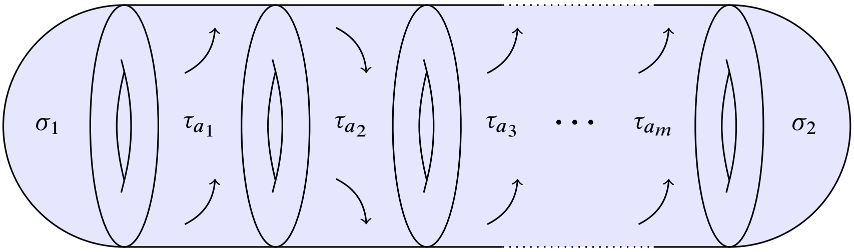

Figure 1: A lens space $L(n,k)$ as solid tori with triangulations σ_1 and σ_2, connected by twisted bundles $(S^1 \times S^1) \rtimes I$ with triangulations τ_{a_j}

that we construct are easy to identify given their specific triangulations. However, the triangulations can then be obfuscated with local moves (eg Newman–Pachner bistellar moves). Proposition 2.1 makes both Theorem 1.1 and the Lackenby–Schleimer result look more interesting. For the latter question, it is easy to compute whether $H_1(M) \cong \mathbb{Z}/n$ is cyclic. If it is, and if n is polynomially bounded in t, then Schleimer's prior result [19] gives an algorithm in NP to compute whether the abelian cover $\widetilde{M}$ is homeomorphic to S^3, which then implies that M is a lens space.

Proposition 2.1 *There exists a family of triangulated lens spaces $\{M \cong L(n,k)\}$ with $t = t(n,k)$ tetrahedra, such that n is exponential in t and there are exponentially many choices for k for each fixed n.*

Proof Our construction is equivalent to a well-known construction of lens spaces using Dehn surgery on a chain of unknots; see Rolfsen [18, Exercise 9H13].

We choose a fixed triangulation σ of the torus $T = S^1 \times S^1$, and we choose two solid tori X_1, X_2 with $\partial X_1 = \partial X_2 = T$, and with triangulations σ_1, σ_2 that extend σ. We can describe an element of the mapping class group of T by an element of $\mathrm{GL}(2,\mathbb{Z})$ that describes its action on the homology group $H_1(T)$. For each $1 \le a \le 5$, we choose a fixed triangulation τ_a of a torus bundle over an interval, $T \rtimes I$, that connects the triangulation σ of T to itself using the monodromy matrix

$$F_a = \begin{pmatrix} 0 & 1 \\ 1 & a \end{pmatrix}.$$

Our construction is to concatenate a sequence $\{\tau_{a_j}\}_{1 \le j \le m}$ of these mapping cylinders together with a solid torus at each end, as in Figure 1. We also assume that $a_1 > 1$. The tetrahedron number t is thus $O(m)$. If the solid tori σ_1 and σ_2 are positioned suitably, then the result is $M \cong L(n,k)$, where n and k are given as a finite continued fraction:

$$\frac{n}{k} = a_m + \cfrac{1}{a_{m-1} + \cfrac{1}{\ddots + \cfrac{1}{a_1}}}.$$

If we let n_j/k_j be the j^{th} partial evaluation, then we can also express the calculation with the recurrence

$$k_j = n_{j-1}, \quad n_j = a_j n_{j-1} + k_j = a_j n_{j-1} + n_{j-2}.$$

The answer n/k determines the monodromy numbers $\{a_j\}$ since the continued fraction is unique under the constraint $a_1 > 1$. Since the integers $\{n_j\}$ increase, we obtain the inequality

$$n_j < (a_j + 1)n_{j-1}.$$

If we choose the sequence of monodromy numbers at random, we obtain the probabilistic relation

$$\mathrm{Ex}[\log n_j] < \mathrm{Ex}[\log(a_j + 1)] + \mathrm{Ex}[\log n_j].$$

Also,

$$\mathrm{Ex}[\log(a_j + 1)] = \tfrac{1}{5}(\log 2 + \log 3 + \cdots + \log 6) < \log 3.73.$$

By the law of large numbers, most monodromy sequences produce $n < 3.73^m$. On the other hand, there are $4 \cdot 5^{m-1}$ sequences of length m, so by the pigeonhole principle, some value of n must see exponentially many values of k. Any such value of n must also be exponentially large. In any case, for every choice of numbers $\{a_j\}$, $\{n_j\}$ grows at least as fast as the Fibonacci numbers, which also implies that n is exponentially large. □

3 Reidemeister torsion

We review Reidemeister torsion (see Turaev [22]) and its value for lens spaces.

Suppose that

$$C_* = \{C_k \xrightarrow{\partial} C_{k-1}\}_{0 \le k \le m}$$

is a finite, acyclic chain complex over a field F. (Reidemeister torsion is well defined for a free complex over any commutative ring, but it is easier to discuss algorithms in the field case.) Suppose in addition that each term C_k has a distinguished basis. Since C_* is acyclic and finite, it is isomorphic to a direct sum of complexes of the form

$$0 \to F \xrightarrow{\cong} F \to 0.$$

Define an *adapted basis* A_* for C_* to be one induced by such a decomposition. In other words, if $\alpha_k \in A_k$ is a basis vector, then either $\partial\alpha_k = 0$ or $\partial\alpha_k \in A_{k-1}$ is another basis vector. Then the *Reidemeister torsion* of C_* is

$$\Delta(C_*) := (\det A_0)(\det A_1)^{-1}(\det A_2)\cdots(\det A_m)^{(-1)^m},$$

where each A_j is also interpreted as the change-of-basis matrix from the distinguished basis to the adapted basis. The following two facts are standard:

(1) Every adapted basis yields the same value of $\Delta(C_*)$.

(2) Let C_* be the chain complex of a finite CW complex σ with PL attaching maps, possibly with twisted coefficients, and using the cells of Φ as its distinguished basis. Then the Reidemeister torsion $\Delta(C_*)$ is invariant under refinement of Φ.

The second fact essentially says that Reidemeister torsion is a PL topological invariant. We have to be careful because the sign of $\Delta(C_*)$ depends on the ordering and orientation of the cells of Φ, and ambiguities in the local coefficient system can also make $\Delta(C_*)$ multivalued.

Let M be a closed, oriented rational homology 3–sphere with a triangulation, or more generally a cellulation which may support a combinatorial local system. We first calculate its untwisted Reidemeister torsion with coefficients in $F = \mathbb{Q}$. Using the orientation, we can canonically augment the chain complex $C_*(M;\mathbb{Q})$ at both ends to obtain the acyclic complex

$$(1)\quad Q_* = \big\{0 \to \mathbb{Q} \to C_3(M;\mathbb{Q}) \to C_2(M;\mathbb{Q}) \to C_1(M;\mathbb{Q}) \to C_0(M;\mathbb{Q}) \to \mathbb{Q} \to 0\big\}.$$

Then it is standard that

$$\Delta(Q_*) = \pm|H_1(M;\mathbb{Z})|.$$

The sign is not a topological invariant, because the j–simplices of M are unordered and unoriented, so they only provide $C_j(M;\mathbb{Q})$ with an unordered, unsigned basis. We choose an ordering and an orientation of the cells such that $\Delta(Q_*) > 0$. We can then use the same ordering and orientation for a Reidemeister torsion calculation on M with twisted coefficients.

Suppose further that

$$H_1(M) = H_1(M;\mathbb{Z}) \cong \mathbb{Z}/n.$$

Then to compute the Reidemeister torsion of M, we let $F = \mathbb{Q}(\zeta_n)$, where ζ_n is an abstract primitive n^{th} root of unity, ie an abstract root of the n^{th} cyclotomic polynomial. We also choose a cellular cocycle $\omega \in C^1(M;\mathbb{Z}/n)$ such that $[\omega]$ generates $H^1(M;\mathbb{Z}/n)$. We use ω to define a twisted coefficient system $\mathbb{Q}(\zeta)_\omega$ on M, and we let

$$R_* := C_*(M;\mathbb{Q}(\zeta_n)_\omega)$$

to define the Reidemeister torsion $\Delta(R_*)$ of M. A change in the choice of the generator $[\omega]$ can change $\Delta(R_*)$ by a Galois automorphism of $\mathbb{Q}(\zeta_n)$. After fixing $[\omega]$, a change in the choice of its representative ω can change $\Delta(R_*)$ by a factor of ζ_n^c for some residue $c \in \mathbb{Z}/n$. Otherwise $\Delta(R_*)$ is a topological invariant of M, provided that the cells of M are ordered and oriented so that $\Delta(Q_*) > 0$.

In particular, if $M = L(n,k)$, then

$$\Delta(R_*) = \zeta_n^c(1-\zeta_n^a)(1-\zeta_n^b), \tag{2}$$

where

$$\frac{a}{b} = k^{\pm 1} \in \mathbb{Z}/n.$$

This answer is easy to calculate using the standard cellulation of $L(n,k)$ with one cell in each dimension, as follows. For a convenient choice of twisted coefficients, this CW complex yields

$$0 \to \mathbb{Q}(\zeta_n) \xrightarrow{1-\zeta_n} \mathbb{Q}(\zeta_n) \xrightarrow{0} \mathbb{Q}(\zeta_n) \xrightarrow{1-\zeta_n^k} \mathbb{Q}(\zeta_n) \to 0.$$

Thus,

$$\Delta(R_*) = (1-\zeta_n)(1-\zeta_n^k).$$

The formula (2) is the same as this one, except generalized to let $\Delta(R_*)$ change with a change in the choice of ω. The exponents a and b are also ambiguous, as follows. First, the value of the torsion (2) does not determine the sign of either a or b, only their relative sign, since

$$\zeta_n^c(1-\zeta_n^a)(1-\zeta_n^b) = \zeta_n^{a+b+c}(1-\zeta_n^{-a})(1-\zeta_n^{-b}).$$

The formula is also symmetric in a and b, so we cannot distinguish k from $1/k$. This stands to reason because

$$L(n,k) \cong L(n,1/k).$$

4 Proof of Theorem 1.1

To prove Theorem 1.1, we begin with two basic results in numerical algorithms.

Theorem 4.1 (Edmonds [3]) *The determinant* $\det M$ *of a square matrix* M *defined over* $\mathbb{Q}(i)$*, the field of complex numbers with rational real and imaginary parts, can be computed in deterministic polynomial time in the bit complexity of* M*.*

Edmonds states his result over an integral domain with suitable arithmetic algorithms; the context of the paper suggests integer matrices. However, his construction works just as well using exact arithmetic in the field $\mathbb{Q}(i)$. He defines a variation of Gaussian elimination such that every number that ever appears is a minor of the original matrix M. As a result, all numbers that arise in the calculation have polynomial bit complexity.

Remark There are many ways to prove Theorem 4.1 and we do not know the best attribution. The hard part of the result is to bound the bit complexity of intermediate expressions, rather than just the number of arithmetic operations.

Theorem 4.1 is related to the problem of calculating the Smith normal form of a matrix.

Theorem 4.2 (Kannan and Bachem [6]) *The Smith normal form of a square or rectangular matrix* M *defined over* $\mathbb{Z}$*, together with left and right multipliers, can be computed in deterministic polynomial time in the bit complexity of* M*.*

Let M be an oriented rational homology 3–sphere described by a triangulation Θ. As a first step which will be important later, we can simplify Θ to a cellulation Φ with one vertex by removing enough triangles until all of the tetrahedra merge into a single 3–cell, and dually by collapsing edges that connect two distinct vertices until only one vertex is left. Using either Θ or Φ, we can calculate the cellular chain complex $C_*(M;\mathbb{Z})$ in polynomial time. We can use Theorem 4.1 to calculate the torsion $\Delta(Q_*)$ of the augmentation Q_* in (1); in particular to determine whether $C_*(M;\mathbb{Z})$ has a positive or negative basis. We can assume a positive basis.

We can iteratively use Theorem 4.2 to calculate a change of basis of the chain complex $C_*(M;\mathbb{Z})$ to put every differential ∂_k into Smith normal form. This also puts the dual complex $C^*(M;\mathbb{Z})$ into Smith normal form. Using Smith normal form, if $H_1(M;\mathbb{Z}) \cong \mathbb{Z}/n$, then we can calculate a cocycle $\omega \in C^1(M;\mathbb{Z}/n)$ that generates $H^1(M;\mathbb{Z}/n)$, and we can express ω in the original basis of $C^*(M;\mathbb{Z}/n)$.

After calculating ω, we can form the chain complex R_* described in Section 3. However, we will want to generalize the calculation, and instead of computing torsion over the abstract field $\mathbb{Q}(\zeta_n)$ which may have exponential dimension over $\mathbb{Q}$, we will compute it over the complex numbers $\mathbb{C}$. To this end, let $\zeta_n = \exp(2\pi i/n)$, and let $\zeta = \zeta_n^\ell$ for certain exponents $\ell \in (\mathbb{Z}/n)^*$. Note that ℓ need not be a prime residue, only nonzero, so ζ may have some lower order $m \mid n$ with $m > 1$. Then we can form the chain complex $R_*(\zeta)$, and its torsion has the same form as in (2):

$$\Delta(R_*(\zeta)) = \zeta^c(1-\zeta^a)(1-\zeta^b). \tag{3}$$

Note that the constants a, b and c depend only on ω and not on the exponent ℓ.

If the cell complex Φ has g edges, then it also has g 2–cells, and we can write the complex $R_*(\zeta)$ as

$$0 \to \mathbb{C} \xrightarrow{\partial_3} \mathbb{C}^g \xrightarrow{\partial_2} \mathbb{C}^g \xrightarrow{\partial_1} \mathbb{C} \to 0.$$

The complex $R_*(\zeta)$ is acyclic, so ∂_3 is injective while ∂_1 is surjective. We can now make an adapted basis as follows:

(1) We use the canonical basis vector $1 \in \mathbb{C}$ in degree 3 of the chain complex $R_*(\zeta)$, and its image under ∂_3 in degree 2. By abuse of notation, we identify the vector $\partial_3(1)$ with the matrix ∂_3, which has only one column.

(2) We choose a nonzero entry of the vector ∂_3. If we choose the j^{th} entry $(\partial_3)_j$, then we can omit the j^{th} canonical basis vector of $\mathbb{C}^g$ in degree 2. We also use the image of these $g-1$ vectors under ∂_2 in degree 1.

(3) We choose a nonzero entry of the dual vector ∂_1. If we choose the k^{th} entry, then we include the k^{th} basis vector of $\mathbb{C}^g$ in degree 1 and its image under ∂_1, which is simply the scalar value $(\partial_1)_k$.

Let $\partial_2^{(j,k)}$ denote the matrix of ∂_2 omitting the j^{th} column and the k^{th} row. Then we can express the Reidemeister torsion of $R_*(\zeta)$ as

$$\Delta(R_*(\zeta)) = \frac{(\partial_3)_j(\partial_1)_k}{\det \partial_2^{(j,k)}}.$$

To compute $\Delta(R_*(\zeta))$ over $\mathbb{C}$, the most important question is how many digits of precision we need throughout the calculation for an accurate final answer.

Lemma 4.3 *Suppose $\zeta = \exp(2\pi i\ell/n) \in \mathbb{C}$ and that $R_*(\zeta)$ is the chain complex of M with its local system $\mathbb{C}_\omega$. Suppose that we want to calculate $z \in \mathbb{Q}(i)$ such that*

$$\Delta(R_*(\zeta)) = z + O(n^{-\alpha})$$

for some constant $\alpha > 0$. Then it suffices to calculate $\det \partial_2^{(j,k)}$ *by estimating its entries with d digits of precision, where d is polynomial in* $\log n$*, g and α. Moreover, the determinant can be calculated in polynomial time.*

Proof Both $(\partial_3)_j$ and $(\partial_1)_k$ are of the form $\zeta^a - \zeta^b$ for some constants a and b, so each of these factors is of order $\Omega(1/n)$. Thus we need to estimate $\det \partial_2^{(j,k)}$ to a precision of $O(n^{-\alpha-2})$. Each entry $\partial_2^{(j,k)}$ is $O(g)$, and therefore each $(g-1)\times(g-1)$ minor of the same matrix is $O(g^{2g})$ since the determinant expansion has $g! = O(g^g)$ terms and each term is $O(g^g)$. So it suffices to estimate each of the $O(g^2)$ terms to precision $O(n^{-\alpha-2}g^{-2g-2})$ in order for $\det \partial_2^{(j,k)}$ (if it is then computed exactly) to have the desired accuracy. Moreover, each term is $O(g)$, which requires $O(\log g)$ digits to the left of each decimal point. Thus the total number of digits need to express each entry is

$$d = O(\log g + \log(n^{\alpha+2}g^{2g+2})),$$

which is polynomial in $\log n$, α, and g. We can then apply Theorem 4.1 to exactly compute the determinant with these approximate entries. □

To complete the proof of Theorem 1.1, recall that $\zeta = \zeta_n^\ell$. Recall from (3) that the Reidemeister torsion is

$$f_-(\zeta) := \Delta(R_*(\zeta)) = \zeta^c(1-\zeta^a)(1-\zeta^b).$$

We want to calculate several values of f to obtained simplified sparse sums:

$$\begin{aligned} f_+(\zeta) &:= \frac{f_-(\zeta^2)}{f_-(\zeta)} = \zeta^c(1+\zeta^a)(1+\zeta^b), \\ g_+(\zeta) &:= \tfrac{1}{2}(f_+(\zeta)+f_-(\zeta)) = \zeta^c + \zeta^{a+b+c}, \\ g_-(\zeta) &:= \tfrac{1}{2}(f_+(\zeta)-f_-(\zeta)) = \zeta^{a+c} + \zeta^{b+c}, \\ h(\zeta) &:= \tfrac{1}{2}(g_+(\zeta)^2 - g_+(\zeta^2)) = \zeta^{a+b+2c}. \end{aligned}$$

At this point we assume that $n > 4$, which we can do since otherwise we can compute the Reidemeister torsion of M directly over the field $\mathbb{Q}(\zeta_n)$. Using the three evaluations

$$\Delta(R_*(\zeta_n)), \quad \Delta(R_*(\zeta_n^2)), \quad \Delta(R_*(\zeta_n^4)), \tag{4}$$

we can learn the sum $g_+(\zeta_n)$ and the product $h(\zeta_n)$ of ζ_n^c and ζ_n^{a+b+c}; and the sum $g_-(\zeta_n)$ and the product $h(\zeta_n)$ of ζ_n^{a+c} and ζ_n^{b+c}. We can thus solve quadratic equations to obtain all four of these numbers. We can then learn the values of an unordered pair $\{\zeta_n^{sa}, \zeta_n^{sb}\}$, where $s = \pm 1$, by taking ratios.

If we calculate the torsion values (4) using floating point arithmetic over $\mathbb{C}$, we obtain floating point approximations to

$$\zeta^{sa} = \exp\Bigl(\frac{2\pi i sa}{n}\Bigr) \quad \text{and} \quad \zeta^{sb} = \exp\Bigl(\frac{2\pi i sb}{n}\Bigr).$$

We can then numerically calculate logarithms to obtain the arguments $2\pi sa/n$ and $2\pi sb/n$. If at this point we know z and $\arg(z)$ to $O(\log n)$ digits of precision, we can calculate the residues $sa, sb \in \mathbb{Z}/n$ by rounding their computed values to the nearest integer. We can then take their ratio in the ring $\mathbb{Z}/n$ to obtain $k^{\pm 1}$. Working backwards, it suffices to compute the values of f_+, $g_\pm$ and h with $O(\log n)$ digits of precision. We can use Lemma 4.3 to specify the precision at the beginning of the calculation in order to have enough precision at this last stage.

References

[1] **S Aaronson**, **C Granade**, **G Kuperberg**, **V Russo**, *The complexity zoo*, electronic resource Available at `http://complexityzoo.com/`

[2] **D Aharonov**, **V Jones**, **Z Landau**, *A polynomial quantum algorithm for approximating the Jones polynomial*, Algorithmica 55 (2009) 395–421 MR

[3] **J Edmonds**, *Systems of distinct representatives and linear algebra*, J. Res. Nat. Bur. Standards Sect. B 71B (1967) 241–245 MR

[4] **S Garnerone**, **A Marzuoli**, **M Rasetti**, *Efficient quantum processing of three-manifold topological invariants*, Adv. Theor. Math. Phys. 13 (2009) 1601–1652 MR

[5] **J Hass**, **J C Lagarias**, **N Pippenger**, *The computational complexity of knot and link problems*, J. ACM 46 (1999) 185–211 MR

[6] **R Kannan**, **A Bachem**, *Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix*, SIAM J. Comput. 8 (1979) 499–507 MR

[7] **G Kuperberg**, *Knottedness is in* NP*, modulo GRH*, Adv. Math. 256 (2014) 493–506 MR

[8] **G Kuperberg**, *Algorithmic homeomorphism of* 3*–manifolds as a corollary of geometrization*, preprint (2015) `arXiv:1508.06720`

[9] **G Kuperberg**, *How hard is it to approximate the Jones polynomial?*, Theory Comput. 11 (2015) 183–219 MR

[10] **G Kuperberg**, *Identifying lens spaces using discrete logarithms, v1*, preprint (2015) `arXiv:1509.02887v1`

[11] **G Kuperberg**, *Learning the exponents in a sum of two modular roots of unity*, MathOverflow post (2015) Available at `https://mathoverflow.net/q/215852`

[12] **M Lackenby**, *The efficient certification of knottedness and Thurston norm*, preprint (2016) `arXiv:1604.00290`

[13] **M Lackenby**, **S Schleimer**, *Lens space recognition is in* NP, Oberwolfach Rep. 9 (2012) 1421–1424 MR

[14] **G Myerson**, *Unsolved problems: how small can a sum of roots of unity be?*, Amer. Math. Monthly 93 (1986) 457–459 MR

[15] **M A Nielsen**, **I L Chuang**, *Quantum computation and quantum information*, Cambridge Univ. Press (2000) MR

[16] **B Poonen**, *Undecidable problems: a sampler*, from "Interpreting Gödel" (J Kennedy, editor), Cambridge Univ. Press (2014) 211–241 MR

[17] **K Reidemeister**, *Homotopieringe und Linsenräume*, Abh. Math. Sem. Univ. Hamburg 11 (1935) 102–109 MR

[18] **D Rolfsen**, *Knots and links*, Mathematics Lecture Series 7, Publish or Perish, Berkeley, CA (1976) MR

[19] **S Schleimer**, *Sphere recognition lies in NP*, from "Low-dimensional and symplectic topology" (M Usher, editor), Proc. Sympos. Pure Math. 82, Amer. Math. Soc., Providence, RI (2011) 183–213 MR

[20] **P W Shor**, *Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer*, SIAM J. Comput. 26 (1997) 1484–1509 MR

[21] **T Tao**, *How small can a sum of a few roots of unity be?*, MathOverflow post (2010) Available at `https://mathoverflow.net/q/46068`

[22] **V Turaev**, *Torsions of* 3*–dimensional manifolds*, Progress in Mathematics 208, Birkhäuser, Basel (2002) MR

Department of Mathematics, University of California
Davis, CA, United States

greg@math.ucdavis.edu

Received: 26 April 2016 Revised: 31 December 2017