AGT

Algebraic & Geometric Topology

Volume 21 (2021)

Coloring invariants of knots and links are often intractable

Greg Kuperberg
Eric Samperton

msp

msp

Algebraic & Geometric Topology 21 (2021) 1479–1510

Coloring invariants of knots and links are often intractable

GREG KUPERBERG
ERIC SAMPERTON

Let G be a nonabelian, simple group with a nontrivial conjugacy class $C \subseteq G$. Let K be a diagram of an oriented knot in S^3, thought of as computational input. We show that for each such G and C, the problem of counting homomorphisms $\pi_1(S^3 \smallsetminus K) \to G$ that send meridians of K to C is almost parsimoniously #P–complete. This work is a sequel to a previous result by the authors that counting homomorphisms from fundamental groups of integer homology 3–spheres to G is almost parsimoniously #P–complete. Where we previously used mapping class groups actions on closed, unmarked surfaces, we now use braid group actions.

20F10, 57M27, 68Q17

1 Introduction

Let K be an oriented knot in the 3–sphere S^3 described by some given knot diagram. Fox [10, Chapter VI, Exercises 6–7] popularized the idea of a 3–coloring of the diagram K, which is now also called a Fox coloring; see de la Harpe and Jones [20]. By definition, such a coloring is an assignment of one of three colors to each arc in K such that at every crossing, the over-arc and the two other arcs are either all the same color or all different colors. It is easy to check that the number of 3–colorings of a diagram is invariant under Reidemeister moves, and is therefore an isotopy invariant of K.

Fox colorings are a special case of the following type of generalized coloring based on the Wirtinger presentation of the knot group $\pi_1(S^3 \smallsetminus K)$ (see Reidemeister [34]): Fix a finite group G and a conjugacy class $C \subseteq G$ that generates G. Then a *C–coloring* is an assignment of an element $c \in C$ to each arc in K such that at each crossing, one of the two relations as in Figure 1 holds, depending on the sign of the crossing. The set of C–colorings is bijective with the set

$$H(K; G, C) := \{f : \pi_1(S^3 \smallsetminus K) \to G \mid f(\gamma) \in C\}$$

Published: 11 August 2021 DOI: 10.2140/agt.2021.21.1479

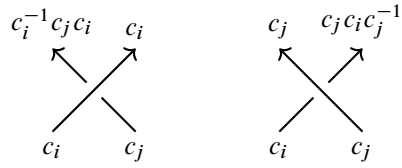

Figure 1: The Wirtinger relations.

of homomorphisms from the knot group to G that take a meridian γ of K to some element in C. (Since the meridians are themselves a conjugacy class of $\pi_1(S^3 \smallsetminus K)$, it doesn't matter which one we choose.) Then $\#H(K;G,C) = |H(K;G,C)|$ is an important integer-valued invariant of knots.

If $G = S_3$ is the symmetric group on three letters and C is the conjugacy class of transpositions, then $H(K;G,C)$ is precisely the set of Fox colorings of K. In this case, and in any case when G is metabelian, $H(K;G,C)$ is an abelian group (or more precisely a torsor over one) that can be calculated efficiently using the Alexander polynomial of K. However, Fox also considered the set $H(K;G,C)$ for general G and C. When $G = A_5$, he observed that "A_5 is a simple group, so that I know of no method of finding representations on A_5 other than just trying" [15]. Our main result, Theorem 1.1, demonstrates that Fox's frustration was prescient, but see Section 1.1 for a more careful discussion.

To state our precise result, we first refine the invariants $H(K;G,C)$ and $\#H(K;G,C)$. Let $\mathrm{Aut}(G,C)$ be the group of automorphisms of G that take C to itself. Then $\mathrm{Aut}(G,C)$ acts on $H(K;G,C)$, and in particular it acts freely on the surjective maps in $H(K;G,C)$. Let

$$Q(K;G,C) := \{f : \pi_1(S^3 \smallsetminus K) \twoheadrightarrow G \mid f(\gamma) \in C\}/\mathrm{Aut}(G,C)$$

be the corresponding quotient set. Regardless of K, the set $H(K;G,C)$ always contains a unique homomorphism with cyclic image that sends γ to each given $c \in C$. If G is not cyclic and if all other homomorphisms are surjective, then, in these cases,

$$\#H(K;G,C) = \#C + \#\mathrm{Aut}(G,C) \cdot \#Q(K;G,C). \tag{1}$$

Our main theorem implies that if G is nonabelian simple, then $\#Q(K;G,C)$ is computationally intractable, and remains so even when every homomorphism $f \in H(K;G,C)$ is promised to be either surjective or have cyclic image.

Theorem 1.1 *Let G be a fixed, finite, nonabelian simple group, and fix a nontrivial conjugacy class $C \subseteq G$. If $K \subseteq S^3$ is an oriented knot specified by a knot diagram*

interpreted as computational input, then the invariant $\#Q(K;G,C)$ is parsimoniously $\#\mathsf{P}$–complete. The reduction also guarantees that $\#Q(K;J,E)=0$ for any group J generated by a conjugacy class E with $\#E<\#C$, except when J is a cyclic group.

We note that J in the statement of the theorem is not necessarily a subgroup of G, although the case that J is a subgroup of G generated by a subset of C is of particular interest.

Before reviewing the definition of $\#\mathsf{P}$–completeness and interpreting Theorem 1.1, we expand on the relation between $H(K;G,C)$ and $Q(K;G,C)$.

Let $c\in C$ and let

$$H(K,\gamma;G,c):=\{f:\pi_1(S^3\smallsetminus K)\to G \mid f(\gamma)=c\}.$$

It is easy to see (by conjugation in G) that $\#H(K,\gamma;G,c)$ is independent of the choice of c and that

$$\#H(K;G,C)=\#C\cdot\#H(K,\gamma;G,c).$$

Let $\mathrm{Aut}(G,c)$ be the group of automorphisms of G that fix c. Then $\mathrm{Aut}(G,c)$ acts on $H(K,\gamma;G,c)$, and in particular it acts freely on the surjective maps in $H(K,\gamma;G,c)$. Let

$$Q(K,\gamma;G,c):=\{f:\pi_1(S^3\smallsetminus K)\twoheadrightarrow G \mid f(\gamma)=c\}/\mathrm{Aut}(G,c)$$

be the corresponding quotient set. Again by examining conjugation in G, we learn that the natural map $Q(K,\gamma;G,c)$ to $Q(K;G,C)$ is a bijection.

Given that every $f\in H(K;G,c)$ has some image $J\ni c$, we obtain the summation formula

$$\#H(K,\gamma;G,c)=\sum_{c\in J\le G}\#\mathrm{Aut}(J,c)\cdot\#Q(K,\gamma;J,c).$$

Given that the conjugacy class of γ generates $\pi_1(S^3\smallsetminus K)$, the conjugacy class E of c in J generates J as well. So we can also write

$$\#H(K;G,C)=\sum_{c\in E\subseteq J\le G}\#C\cdot\#\mathrm{Aut}(J,E)\cdot\#Q(K;J,E). \tag{2}$$

Finally, if $J\neq G$, then necessarily $\#E<\#C$. Thus when the conclusion of Theorem 1.1 holds, equation (2) reduces to equation (1).

1.1 Interpretation and previous results

For an introduction to the topic of computational complexity, see our previous article [28, Section 2.1], as well as Arora and Barak [3] and *The complexity zoo* [1]. Here

we just give a brief description the concept of #P–completeness and parsimonious reduction.

If A is a finite alphabet and A^* is the set of finite words in A, a problem in #P is by definition a function $c\colon A^* \to \mathbb{N}$ given by the equation

$$c(x) = \#\{y \mid p(x, y) = \text{yes}\},$$

where the length of the certificate y is polynomial in the length of x and

$$p\colon A^* \times A^* \to \{\text{yes}, \text{no}\}$$

is a predicate that can be computed in polynomial time. A counting problem $c \in \#\mathsf{P}$ is *parsimoniously* #P*–complete* when every problem $b \in \#\mathsf{P}$ can be converted to a special case of c. More precisely, c is parsimoniously #P–complete when $b(x) = c(f(x))$ for some function $f\colon A^* \to A^*$ that can be computed in polynomial time.

The significance of parsimonious #P–completeness for a counting problem c is that not only is the exact value of c computationally intractable, but also obtaining any partial information about c is computationally intractable, assuming standard conjectures in complexity theory. To give a contrasting example, the number of perfect matchings $m(\Gamma)$ of a finite, bipartite graph Γ is well known to be #P–complete by the looser standard of Turing–Cook reduction [39]. The exact value of $m(\Gamma)$ is thus intractable. However, the parity of m can be computed in polynomial time (as a determinant over $\mathbb{Z}/2$), whether $m(\Gamma)$ equals 0 can be computed in polynomial time (see Munkres [33]), and $m(\Gamma)$ can be approximated in randomized polynomial time (see Jerrum, Sinclair and Vigoda [24]). Barring a catastrophe in computer science, no such partial results are possible for computing $\#Q(K; G, C)$ under the hypotheses of Theorem 1.1, not even with the aid of a quantum computer; see Bennett, Bernstein, Brassard and Vazirani [5].

The analogous concepts for existence questions are the complexity class NP and the NP–completeness property. A decision function $d\colon A^* \to \{\text{yes}, \text{no}\}$ is in NP if there is a polynomial-time predicate p such that $d(x) = \text{yes}$ if and only if $p(x, y) = \text{yes}$. The function d is *Post–Karp* NP*–complete* if for every $e \in \mathsf{NP}$, $e(x) = d(f(x))$ for some f computable in polynomial time.

In particular, Theorem 1.1 implies that the existence of a nontrivial C–coloring of a knot is NP–complete. In fact, Theorem 1.1 implies much more thanks to a reduction due to Valiant and Vazirani [40]. Namely, distinguishing *any* two values of a parsimoniously #P–hard problem is NP–hard with randomized reduction [28, Theorem 2.1].

Some partial information about the unadjusted counting invariant $\#H(K;G,C)$ can be computed efficiently, for instance that it is always at least $\#C$. However, Theorem 1.1 and (1) together imply that this extra information can be trivial. We call a counting problem $c \in \#\mathsf{P}$ *almost parsimoniously* $\#\mathsf{P}$*–complete* if for every $b \in \#\mathsf{P}$, there is a reduction $\alpha b(x)+\beta = c(f(x))$ for some universal constants $\alpha > 0$ and $\beta \geq 0$. Almost parsimonious reductions arise naturally in computational complexity. For example, the number of 3–colorings of a planar graph with at least one edge is always divisible by 6; but, after dividing by 6, this number becomes parsimoniously $\#\mathsf{P}$–complete; see Barbanchon [4]. Likewise, Theorem 1.1 shows $\#H(K;G,C)$ is almost parsimoniously $\#\mathsf{P}$–complete.

De Mesmay, Rieck, Sedgwick and Tancer [32] state that, besides their results and ours, they know of no other hardness results concerning knots in the 3–sphere (as opposed to links in the 3–sphere, by Lackenby [29] and Koenig and Tsvietkova [25]; or knots in other 3–manifolds, by Agol, Hass and Thurston [2]). In fact, the first author previously showed that it is $\#\mathsf{P}$–hard to approximate the Jones polynomial of a knot [27], using programmability methods that originate with quantum computing; see Bordewich, Freedman, Lovász and Welsh [7]. As in earlier work of Krovi and Russell [26], here we use an analogous classical programmability approach for classical knot invariants. We should also mention the prescient result of Jaeger that exact computation of the HOMFLY polynomial is $\#\mathsf{P}$–complete [22], which was soon followed by the same result for the Jones polynomial, due to Jaeger, Vertigan and Welsh [23]. These early hardness results were for links rather than knots, but it was widely assumed (albeit without further investigation until recently) that they probably also hold for knots.

The strongest partial result toward Theorem 1.1 to our knowledge is that of Krovi and Russell. Taking the straightforward generalization of $H(K;G,C)$ to links L, they showed that $\#H(L;A_m,C)$ is $\#\mathsf{P}$–complete for any fixed $m \geq 5$ and any fixed conjugacy class C of permutations with at least four fixed points. Their reduction is not almost parsimonious because it has an error term. In particular, they do not obtain that it is NP–complete to determine whether $\#H(L;A_m,C) > \#C$ or $\#Q(L;A_m,C) > 0$.

We conclude with some comments, partly due to one of the referees, on the refined complexity of coloring invariants of knot diagrams with n crossings. First, there is a better way to compute $\#Q(K;G,C)$ than "just trying" if we take that to be an exhaustive search over $2^{O(n)}$ colorings of the edges to see which ones are admissible. A theorem of Lipton and Tarjan [30] implies that every planar graph Γ has pathwidth $O(\sqrt{n})$, and there is a polynomial-time algorithm to find a position for Γ with this width

[6, Corollary 23]. Consequently, a standard dynamic programming algorithm (which may more generally use treewidth rather than pathwidth) can calculate $\#H(K;G,C)$ in time $2^{\widetilde{O}(\sqrt{n})}$. We can then use Möbius inversion over subgroups of G to calculate $\#Q(K;G,C)$. Meanwhile, the reductions from #CSAT to #ZSAT in our proof of Theorem 1.1 are all linear expansions for any fixed G, with the exception of a quadratic expansion to make circuits planar [28, Section 2.2]. The exponential-time hypothesis (ETH) of Impagliazzo and Paturi [21], together with the proof of Theorem 1.1, thus implies a time complexity lower bound of $2^{\widetilde{\Omega}(\sqrt{n})}$, as has previously been noted for related problems for planar graphs; see Lokshtanov, Marx and Saurabh [31]. Thus Theorem 1.1 is sharp under ETH, up to a constant factor α_G in the exponent that depends on G. However, our methods produce poor lower bounds for α_G.

1.2 Outline of the proof

Our proof of Theorem 1.1 follows our proof of the analogous theorem for homology 3–spheres [28], which we assume as a prerequisite for this article. However, Theorem 1.1 is a stronger result because knots are a more restricted class of topological objects. As a preliminary observation, both $\#H(K;G,C)$ and $\#Q(K;G,C)$ are in #P by the same argument as in the 3–manifold case [28, Theorem 2.7].

The reduction begins with a counting version of circuit satisfiability, #CSAT, that is rather directly parsimoniously #P–complete [28, Theorem 2.2]. The #CSAT problem can be reduced to a certain version with reversible circuits, #RSAT, and we can assume in both problems that circuits are planar. Whereas the output to a CSAT circuit is constrained to yes and the input is any satisfying certificate, both the input and output of a #RSAT circuit are partially constrained. In turn, #RSAT reduces almost parsimoniously to an ad hoc reversible circuit problem, called #ZSAT, where

(1) the alphabet is a U–set for some finite group U with a single fixed point called the "zombie" symbol and otherwise free orbits, and

(2) the gates are U–equivariant permutations.

Finally, #ZSAT reduces to $\#Q(K;G,C)$ in a construction in which the circuit becomes a braid word and suitable initialization and finalization conditions are expressed by plat closure.

Let $D^2\smallsetminus[n]$ denote a disk with n punctures. The reduction from #ZSAT to $\#Q(K;G,C)$ involves a braid group action on the set of surjections

$$f:\pi_1(D^2\smallsetminus[2k])\twoheadrightarrow G$$

with clockwise monodromy in C at k punctures, counterclockwise monodromy in C at the other k punctures, and trivial monodromy on the outside. In Theorem 4.7, we show that when k is large enough, this braid group action is very highly transitive modulo a certain Schur invariant. High transitivity makes it possible to implement gates in a precise way that preserves enumeration and does not disturb nonsurjective homomorphisms. Theorem 4.7 in turn requires two types of group-theoretic ingredients. The first ingredient, Theorem 4.2, is a refinement of the Conway–Parker theorem (see Fried and Völklein [16]) that shows that the action is at least transitive when G is any finite group. This refinement first appeared in a retracted e-print of Ellenberg, Venkatesh and Westerland [13, Theorem 7.6.1]; the second author of this article later found a topological proof [38, Theorem 1.1]. The second ingredient is a set of surjectivity results for group homomorphisms (Section 2).

Acknowledgement The authors were partly supported by NSF grant CCF-1716990.

2 Group theory

In this section, we simply list some surjectivity results in group theory that we will need to prove Theorem 4.7.

2.1 Surjectivity for products

The following lemma is a mutual corollary of Goursat's lemma [18] and Ribet's lemma [36; 35]. In our research, we first saw it stated by Dunfield and Thurston [12, Lemma 3.7].

Lemma 2.1 (after Goursat and Ribet [36, Lemmas 5.2.1 and 5.2.2]) *If*

$$f : J \to G_1 \times G_2 \times \cdots \times G_n$$

is a homomorphism from a group J to a product of nonabelian simple groups that surjects onto each factor, and if no two factor homomorphisms $f_i : J \twoheadrightarrow G_i$ and $f_j : J \twoheadrightarrow G_j$ are equivalent by an isomorphism $G_i \cong G_j$, then f is surjective.

Remark Results similar to Lemma 2.1 have appeared many times in the literature with various attributions and extra hypotheses. Both Ribet and Dunfield and Thurston assume that the target groups are finite even though their proofs do not use this hypothesis.

Dunfield and Thurston also state that the result is due to Hall [19]. However, all that we can find in this citation is an unproven lemma (in his Section 1.6) that can (with a little work) be restated as a special case of Lemma 2.1.

Say that a group G is *Zornian* if every proper normal subgroup of G is contained in a maximal normal subgroup. Clearly every finite group is Zornian, which is the case that we will need; more generally, every finitely generated group is Zornian [28, Section 3.2]. The following is also an adaptation of Goursat's lemma.

Lemma 2.2 (after Goursat [28, Lemma 3.6]) *Suppose that*

$$f : J \to G_1 \times G_2$$

is a group homomorphism that surjects onto G_1, and suppose that G_1 is Zornian. If no simple quotient of G_1 is involved in G_2, then $f(J)$ contains G_1.

Finally, we will need the following two related lemmas.

Lemma 2.3 (Ribet [36, Section 5.2]) *If*

$$N \trianglelefteq G = G_1 \times G_2 \times \cdots \times G_n$$

is a normal subgroup of a product of perfect groups that surjects onto each factor G_i, then $N = G$.

Lemma 2.4 [28, Lemma 3.7] *If*

$$f : G_1 \times G_2 \times \cdots \times G_n \twoheadrightarrow J$$

is a surjective homomorphism from a direct product of groups to a nonabelian simple quotient J, then it factors through a quotient map $f_i : G_i \twoheadrightarrow J$ for a single value of i.

2.2 Rubik groups

If X is a finite set, then $\mathrm{Sym}(X)$ and $\mathrm{Alt}(X)$ denote the symmetric and alternating groups on X, while $\mathrm{Sym}(n) = S_n$ and $\mathrm{Alt}(n) = A_n$ for any positive integer n. If G is a group and X is a G–set with finitely many orbits, then we denote the group of G–set automorphisms of X by $\mathrm{Sym}_G(X)$. We define the *Rubik group* to be the commutator subgroup

$$\mathrm{Rub}_G(X) := [\mathrm{Sym}_G(X), \mathrm{Sym}_G(X)].$$

Note that the natural map $\mathrm{Sym}_G(X) \to \mathrm{Sym}(X/G)$ takes $\mathrm{Rub}_G(X)$ to $\mathrm{Alt}(X/G)$.

When X is a free G–set with $\#(G/X) = n$, $\mathrm{Sym}_G(X)$ is isomorphic to the restricted wreath product

$$\mathrm{Sym}(n, G) := G \operatorname{wr}_m \mathrm{Sym}(n) = G^{\times n} \rtimes \mathrm{Sym}(n).$$

Likewise, let

$$\mathrm{Rub}(n, G) := [\mathrm{Sym}(n, G), \mathrm{Sym}(n, G)].$$

We need two results about Rubik groups from our previous work [28], which we will restate here.

The first result is an elementary counterpart for Rubik groups to the well-known corollary of the classification of finite simple groups that a 6–transitive subgroup of $\mathrm{Sym}(n)$ is *ultratransitive*, meaning that it contains $\mathrm{Alt}(n)$ or equivalently that it is $(n-2)$–transitive. Say that a group homomorphism

$$f : J \to \mathrm{Sym}(n, G)$$

is *G–set k–transitive* if it acts transitively on ordered lists of k elements that all lie in distinct G–orbits. Say likewise that it is *G–set ultratransitive* if its image contains $\mathrm{Rub}(n, G)$.

Theorem 2.5 [28, Theorem 3.10] *Let G be a group and let $n \geq 7$ be an integer such that $\mathrm{Alt}(n-2)$ is not a quotient of G. Suppose that a homomorphism*

$$f : J \to \mathrm{Sym}(n, G)$$

from a group J is G–set 2–transitive and that its projection $\mathrm{Rub}(n, G) \to \mathrm{Alt}(n)$ is 6–transitive (and therefore ultratransitive). Then f is G–set ultratransitive.

The other result says $\mathrm{Rub}(n, G)$ has a unique simple quotient when $\mathrm{Alt}(n)$ is a simple group.

Lemma 2.6 [28, Lemma 3.11] *If G is any group and $n \geq 5$, then the only simple quotient of $\mathrm{Rub}(n, G)$ is $\mathrm{Alt}(n)$.*

3 Equivariant circuits

In this section, we review the ZSAT circuit model from our previous work [28].

Let A be a finite set with at least two elements, considered as a computational alphabet. A *reversible circuit of width n* is a bijection $A^n \to A^n$ expressed as a composition of

bijective gates $A^k \to A^k$ in the pattern of a directed, acyclic graph. The gates are all chosen from some fixed finite set of bijections.

Let G be a nontrivial finite group acting on A with a single fixed point z, the *zombie symbol*, and otherwise with free orbits. Let I and F be two proper G–invariant subsets of A that contain the zombie symbol and are not just that symbol:

$$\{z\} \subsetneq I, F \subsetneq A.$$

We interpret I as an initial subalphabet and F as a final subalphabet. An instance Z of $\mathsf{ZSAT}_{G,A,I,F}$ is a planar reversible circuit with gate set $\mathrm{Rub}_G(A^2)$. (This gate set then generates $\mathrm{Rub}_G(A^k)$ for each $k > 2$.) Then a certificate accepted by Z is a solution to the constraint satisfaction problem

$$x \in I^n \quad \text{and} \quad Z(x) \in F^n,$$

where n is the width of Z. The counting problem $\#\mathsf{ZSAT}_{G,A,I,F}$ counts the number of such solutions to Z.

We will need the following technical result from our previous work:

Theorem 3.1 [28, Lemma 4.1] *As above, let G act on the alphabet A with a single fixed point z, and otherwise freely, and let I and F be two proper G–invariant subsets with $\{z\} \subsetneq I$, $F \subsetneq A$. Then $\#\mathsf{ZSAT}_{G,A,I,F}$ is almost parsimoniously $\#\mathsf{P}$–complete. If p is a counting problem in $\#\mathsf{P}$, then there is a polynomial-time reduction $f \in \mathsf{FP}$ such that*

$$\#\mathsf{ZSAT}_{G,A,I,F}(f(x)) = \#G \cdot p(x) + 1$$

for every instance x of p, where the $+1$ term accounts for the trivial, all zombie solution $(z, \dots, z)$. More precisely, the number of free orbits of nontrivial solutions is parsimoniously $\#\mathsf{P}$–complete.

Remark In our previous work, we did not put the zombie symbol z in the sets I and F, instead setting the initial and final sets to be $(I \cup \{z\})^n$ and $(F \cup \{z\})^n$. We also assumed the side conditions that $I \neq F$, that $\#A \geq 2\,\#(I \cup F) + 3\,\#G + 1$, and that $\#I, \#F \geq 2\,\#G$. The first two of these conditions were recognized as optional, but in fact they are all optional. We can emulate the first condition by adding a layer of unary gates at the beginning or end, and we can attain the other two conditions by replacing (A, I, F) by (A^k, I^k, F^k) for some constant k.

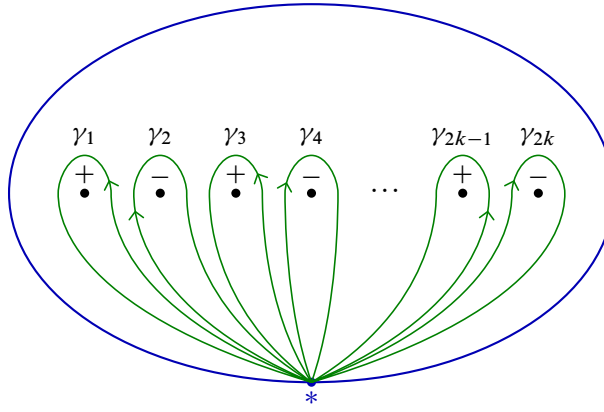

Figure 2: Generators for $\pi_1(D^2 \smallsetminus [2k])$.

4 Braid group actions

The main goal of this section is Theorem 4.7. This theorem is a refinement, in the special case that G is simple, of a result of Roberts and Venkatesh [37, Theorem 5.1].

4.1 Conjugacy-restricted homomorphisms and colored braid subgroups

Recall that $D^2 \smallsetminus [2k]$ denotes the disk D^2 minus a set $[2k]$ of $2k$ points. As shown in Figure 2, place the points in a line and alternately label them $+$ and $-$, and choose a basepoint $*$ that is not on the same line. Also as shown, choose a list of generators of $\pi_1(D^2 \smallsetminus [2k])$ represented by simple closed curves $\gamma_1, \dots, \gamma_{2k}$, where each γ_i winds counterclockwise around the puncture p_i when it is positive (when i is odd) and clockwise when it is negative (when i is even). Finally we choose one more curve

$$\gamma_\infty = \gamma_{2k}^{-1}\gamma_{2k-1}\cdots\gamma_4^{-1}\gamma_3\gamma_2^{-1}\gamma_1$$

representing the boundary of the disk. Note that we concatenate from left to right, eg $\gamma_1\gamma_2$ is the element of $\pi_1(D^2 \smallsetminus [2k])$ that first traverses γ_1 and then γ_2.

When G is a finite group and $C \subset G$ is a union of conjugacy classes, we define three sets of homomorphisms from $\pi_1(D^2 \smallsetminus [2k])$ to G:

$$\begin{aligned} T_k(G,C) &:= \{f : \pi_1(D^2 \smallsetminus [2k]) \to G \mid f(\gamma_i) \in C\}, \\ \widehat{R}_k(G,C) &:= \{f \in T_k(G,C) \mid f(\gamma_\infty) = 1\}, \\ R_k(G,C) &:= \{f \in \widehat{R}_k(G,C) \mid f \text{ is onto}\}. \end{aligned}$$

When G and C are clear from context, we will also use the abbreviations

$$T_k := T_k(G,C), \quad \widehat{R}_k := \widehat{R}_k(G,C), \quad R_k := R_k(G,C).$$

Note that C will often but not always be a single conjugacy class.

Since $\pi_1(D^2 \smallsetminus [2k])$ is freely generated by $\gamma_1, \gamma_2, \dots, \gamma_{2k}$, the set $T_k(G,C)$ is bijective with C^{2k}. By abuse of notation, we will sometimes specify elements of $T_k(G,C)$ as lists of elements in C.

Since $f(\gamma_\infty) = 1$ in both $\widehat{R}_k$ and R_k, each such f factors through $\pi_1(S^2 \smallsetminus [2k])$, where $S^2 \smallsetminus [2k]$ is a punctured sphere obtained by collapsing the boundary ∂D^2 to a point. Also by abuse of notation, we will interpret each such f as having this domain instead.

The homomorphism sets $T_k(G,C)$, $\widehat{R}_k(G,C)$ and $R_k(G,C)$ are all invariant under the colored braid group $B_{k,k} \leq B_{2k}$, by definition the subgroup of the braid group that preserves the labels $+$ and $-$ of the $2k$ punctures. The goal of this section is to show that the action of $B_{k,k}$ is large enough that we can implement gates in ZSAT with it.

4.2 Invariant homology classes and the Conway–Parker theorem

In this subsection, let G be any finite group which is generated by a single conjugacy class C. We describe the orbits of the $B_{k,k}$–action on R_k in the limit as $k \to \infty$. We will say that a property of the action holds *eventually* if it is a stable property in this limit, ie if it is true for all k large enough.

The main tool that we need is the Brand classifying space $B(G,C)$ [8]. This space is a modification of the usual classifying space $B(G)$ (often written BG) of a group G "relative" to a conjugacy class C. It has a number of important properties for our purposes, some described by Brand, and some given by the second author [38]. Before stating these properties, we first review the definition. The free loop space $LB(G)$ comes with an evaluation map

$$\mathrm{ev}\colon LB(G) \times S^1 \to B(G),$$

and it has a connected component $L_C B(G)$ whose loops represent the chosen conjugacy class $C \subseteq G$. We define $B(G,C)$ by gluing $B(G)$ to $L_C B(G) \times D^2$ using the evaluation map:

$$B(G,C) := (B(G) \sqcup L_C B(G) \times D^2)/\mathrm{ev}. \tag{3}$$

In other words, an element $f \in L_C B(G)$ is also a continuous function $f: S^1 \to B(G)$, and we identify $(f, x) \in L_C B(G) \times D^2$ with $f(x) \in B(G)$ when $x \in S^1 = \partial D^2$. We also retain a basepoint for $B(G)$ and thus $B(G, C)$, even though we use the free loop space rather than the based loop space to define the latter.

Recall also that the homology of a group is by definition the homology of a classifying space, $H_*(G) = H_*(B(G))$.

Remark Following Ellenberg, Venkatesh and Westerland [13], in [37], Roberts and Venkatesh extend the notation $H_*(G)$ in an ad hoc way to a reduced Schur multiplier, which they denote by "$H_2(G, C)$". We will later denote the reduced Schur multiplier as $M(G, C)$ instead. It is a subgroup of $H_2(B(G, C))$, which we will also not abbreviate as "$H_2(G, C)$". The reason is that we do not know a natural interpretation of either $M(G, C)$ or $H_*(B(G, C))$ as a relative homology group.

Recall that if A is an abelian group, then A_{tor} is its torsion subgroup and $A_{\text{free}} = A/A_{\text{tor}}$ is its free quotient.

Proposition 4.1 [38] *Let $B(G, C)$ be the Brand classifying space for a finite group G generated by a conjugacy class C. For each a and b, let $\Sigma_{a,b} = S^2 \smallsetminus [a + b]$ be a punctured sphere with a points marked $+$, b points marked $-$ and a basepoint $*$. Let $[S^2] \in H_2(S^2)$ be a fixed orientation of S^2. Then*:

(1) *Let $f: \pi_1(\Sigma_{a,b}) \to G$ be a group homomorphism such that each $+$ point has counterclockwise monodromy in C and each $-$ point has clockwise monodromy in C. Then f is represented by a pointed map $\phi: S^2 \to B(G, C)$.*

(2) *Every pointed map $\phi: S^2 \to B(G, C)$ in general position yields a homomorphism $f: \pi_1(\Sigma_{a,b}) \to G$ as in (1) for some a and b. The maps $\phi_0 \sim \phi_1$ are homotopic if and only if the homomorphisms f_0 and f_1 are connected by a concordance*

$$f: \pi_1((S^2 \times I) \smallsetminus L) \to G,$$

where L is a tangle.

(3) *Given that C generates G, $B(G, C)$ is simply connected.*

(4) *There is an exact sequence*

$$Z(c)_{\text{ab}} \xrightarrow{\kappa} H_2(G) \xrightarrow{\beta} H_2(B(G, C)) \xrightarrow{\sigma} \mathbb{Z} \to 0,$$

where $Z(c) \subseteq G$ is the centralizer of any one element $c \in C$. Given that G is finite, the image of β is $H_2(B(G, C))_{\text{tor}}$.

(5) *Given* ϕ, *and* σ *from* (4),

$$(\sigma \circ \phi_*)([S^2]) = a - b \in \mathbb{Z}.$$

Expanding on Proposition 4.1(1): For a $+$ point, if $\gamma \in \pi_1(\Sigma_{a,b})$ is freely homotopic to the boundary of a small disk centered at that point and containing no other points, with a parametrization counterclockwise with respect to the orientation $[S^2]$, then $f(\gamma) \in C$. For a $-$ point, $f(\gamma) \in C$ when γ winds *clockwise* with respect to the orientation.

Remark The Brand classifying space $B(G, C)$ exists for any G (not necessarily finite) and any union of conjugacy classes (not necessarily only one, and not necessarily generating G). In general, the homotopy classes $[M, B(G, C)]$ from a smooth manifold M of any dimension classify the concordance classes of C–branched G–covers of M, such that the codimension 2 branch locus has a distinguished normal framing. Likewise, the cobordism groups $\Omega_n(B(G, C))$ classify the cobordism classes of such branched coverings of n–manifolds.

Following Proposition 4.1, we introduce the abuse of notation

$$f_* = \phi_* \colon H_2(S^2) \to H_2(B(G, C)).$$

We also define

$$M(G, C) := H_2(B(G, C))_{\text{tor}}.$$

The group $M(G, C)$, the *reduced Schur multiplier*, is a quotient of the usual Schur multiplier $M(G) = H_2(G)$. Now suppose that

$$f \colon \pi_1(S^2 \smallsetminus [2k]) \to G$$

is an element of $\widehat{R}_k$, with $k = a = b$. Then $a - b = 0$, so by parts (4) and (5) of the proposition,

$$f_*([S^2]) \in H_2(B(G, C))$$

maps to zero in $H_2(B(G, C))_{\text{free}}$. It thus lies in $M(G, C)$. We define the (branched) *Schur invariant* of f to be

$$\text{sch}(f) := f_*([S^2]) \in M(G, C).$$

By construction, the function

$$\text{sch} \colon \widehat{R}_k \to M(G, C)$$

is invariant under the action of the colored braid group $B_{k,k}$.

Theorem 4.2 (Ellenberg, Venkatesh and Westerland [38, Theorem 1.1]) *Let G be a finite group generated by a conjugacy class C. Then, eventually, the Schur invariant yields a bijection*

$$\mathrm{sch}\colon R_k/B_{k,k} \xrightarrow{\cong} M(G,C).$$

In particular, sch *is eventually injective; ie it is eventually a complete orbit invariant for the action of $B_{k,k}$ on R_k.*

Remark Theorem 4.2 first appeared in an arXiv e-print by Ellenberg, Venkatesh and Westerland [13]. This e-print was later withdrawn for unrelated reasons, but (besides that arXiv versions are permanent) the argument was later cited and sketched by Roberts and Venkatesh [37]. The second author [38] then found a topological proof of the same result using the Brand classifying space. The new results of [38] also hold for surfaces with either genus or punctures or both, and thus subsume a result of Dunfield and Thurston [12, Theorem 6.23]. In fact, Theorem 4.2 also holds when C is a union of conjugacy classes rather than just one. (In full generality, the Schur invariant $\mathrm{sch}(f)$ lies in a torsor of the reduced Schur multiplier $M(G,C)$ rather than directly in this abelian group.) The original result along these lines is the unpublished Conway–Parker theorem, which is the case $C = G$, and which was later proven in the literature by Fried and Völklein [16].

We will use two basic properties of the Schur invariant, one of which requires a definition: say that $f \in \widehat{R}_k$ *bounds a plat* if there is an inclusion $S^2 \smallsetminus [2k] \to B^3$ such that f extends to a homomorphism from the fundamental group of the complement of a trivial tangle in B^3 with oriented arcs.

Lemma 4.3 *The Schur invariant has the following properties*:

(1) *If $f \in \widehat{R}_a$ and $g \in \widehat{R}_b$, with boundary sum $f \# g \in \widehat{R}_{a+b}$, then* $\mathrm{sch}(f \# g) = \mathrm{sch}(f) + \mathrm{sch}(g)$.

(2) *If $f \in \widehat{R}_k$ bounds a plat, then* $\mathrm{sch}(f) = 0$.

Proof (1) The boundary sum of f and g corresponds to the group law in

$$\pi_2(B(G,C)) = [S^2, B(G,C)].$$

Thus,

$$\mathrm{sch}(f \# g) = (f \# g)_*([S^2]) = f_*([S^2]) + g_*([S^2]) = \mathrm{sch}(f) + \mathrm{sch}(g).$$

(2) The map f is nullconcordant by hypothesis, hence $\mathrm{sch}(f) = 0$ is nullhomologous by Proposition 4.1. □

Our reduction from ZSAT to $\#H(G,C)$ makes special use of maps $f \in \widehat{R}_k$ with $\mathrm{sch}(f) = 0$. Hence we define

$$\widehat{R}_k^0 := \{f \in \widehat{R}_k \mid \mathrm{sch}(f) = 0\}, \quad R_k^0 := \{f \in R_k \mid \mathrm{sch}(f) = 0\}.$$

4.3 The perfect case

In this subsection, we establish some further properties of the reduced Schur multiplier $M(G,C)$ with the additional assumption that G is perfect. Not all of the properties require this hypothesis, but everything listed is at least better motivated in that case. We begin with the following interpretation of $M(G,C)$, which is explained by Roberts and Venkatesh [37, Section 4B]. If G is a perfect group, then it has a canonical central extension

$$M(G) \hookrightarrow \widehat{G} \twoheadrightarrow G,$$

called the *Schur cover* $\widehat{G}$ of G. In general the conjugacy classes in $\widehat{G}$ can be larger than their counterparts in G, in the sense that two preimages $g_1, g_2 \in \widehat{G}$ of one element $g \in G$ can be conjugate to each other. The reduced multiplier $M(G,C)$ is the finest possible quotient of $M(G)$ such that, in the corresponding central extension

$$M(G,C) \hookrightarrow \widetilde{G} \twoheadrightarrow G,$$

two distinct preimages $c_1, c_2 \in \widetilde{G}$ of $c \in C$ are never conjugate. In other words, if $C' \subseteq \widetilde{C}$ is any one conjugacy class in the preimage $\widetilde{C}$ of C, then $\widetilde{C}$ decomposes as

$$\widetilde{C} = M(G,C) \cdot C',$$

where each mC' with $m \in M(G,C)$ is a distinct conjugacy class.

Lemma 4.4 *If C generates G and G is perfect, then*

$$\lim_{k\to\infty} \frac{\#R_k}{(\#C)^{2k}} = \frac{1}{\#G}, \quad \lim_{k\to\infty} \frac{\#R_k^0}{(\#C)^{2k}} = \frac{1}{\#G \cdot \#M(G,C)}.$$

Dunfield and Thurston [12, Lemmas 6.10 and 6.13] proved a version of Lemma 4.4 for maps from fundamental groups of closed surfaces instead of punctured disks. Roberts and Venkatesh [37, equations (3.7) and (4.6)] proved a version for punctured disks, but with the homomorphism sets quotiented by the conjugation action of G.

Proof We begin by establishing the first limit, except with a numerator of $\#\widehat{R}_k$ rather than $\#R_k$. Consider an infinite list

$$c_1, c_2, c_3, \ldots \in C$$

of elements of C chosen independently and uniformly at random. The first $2k$ of these elements describe a homomorphism

$$f_k : \pi_1(D^2 \smallsetminus [2k]) \to G$$

with $f_k \in T_k$ following the conventions in Section 4.1. Then $f_k \in R_k$ when the product

$$g_k := c_{2k}^{-1} c_{2k-1} \cdots c_3 c_2^{-1} c_1$$

equals 1, since $g_k = f_k(\gamma_\infty)$. The first limit can thus be restated as saying that the probability that $g_k = 1$ converges to $1/\#G$ as $k \to \infty$. To argue this, we note the inductive relation

$$g_k = c_{2k}^{-1} c_{2k-1} g_{k-1},$$

and we let M be the corresponding stochastic transition matrix, independent of k, on probability distributions on g_k drawn from G. One can check these three properties of M:

(1) M commutes with both left and right multiplication by G.

(2) M is symmetric, $M = M^T$, and thus doubly stochastic.

(3) Each diagonal entry of M equals $1/\#C$.

We apply the Perron–Frobenius theorem to the matrix M, in the doubly stochastic case. By this theorem, either M^k converges to a constant matrix as $k \to \infty$, or G has a nontrivial equivalence relation $\sim$ such that M descends to a permutation on the quotient set $G/\!\sim$. Since the diagonal of M is entirely nonzero, this permutation must be the identity. Since M commutes with both left and right multiplication by every $g \in G$, the set quotient $G/\!\sim$ must be a group quotient G/N by some normal group $N \trianglelefteq G$. The conjugacy class C descends to a conjugacy class E which generates G/N. Since M acts by the identity on G/N, E must have a single element. Thus G/N would be a cyclic group if it existed, contradicting that G is perfect.

This establishes the first limit with numerator $\widehat{R}_k$ instead of R_k. For the stated limit, observe that in the given random process, the image of f_k is monotonic; more precisely, it almost surely increases to G and stays there. By contrast, the condition $g_k = 1$ is recurrent. Therefore the limiting probability that $f_k \in \widehat{R}_k$ is the same as the limiting probability that $f_k \in R_k$.

We reduce the second limit to the first one, just as Roberts and Venkatesh reduce their equation (4-6) to their equation (3-7) [37]. Recall that $\widetilde{G}$ is the central extension of G by $M(G, C)$, and that $\widetilde{G}$ is a perfect group because the full Schur cover $\widehat{G}$ of the

perfect group G is perfect. Let $C' \subseteq \widetilde{G}$ be a conjugacy class that lifts C. (Any such lift generates $\widetilde{G}$.) Then $f \in R_k$ has a surjective lift $f' \in T_k(\widetilde{G}, C')$. Following [37, Section 4E], we can recognize the Schur invariant $\mathrm{sch}(f)$ as

$$\mathrm{sch}(f) = f'(\gamma_\infty),$$

independent of the choice of C'. Therefore the second limit for the group G is equivalent to the first limit for the group $\widetilde{G}$, as desired. □

The remaining properties concern the Cartesian power (G^ℓ, C^ℓ) of (G, C) and require some algebraic topology to state properly. If we identify $B(G^\ell)$ with $B(G)^\ell$, then this identification extends to a natural map

$$\psi\colon B(G^\ell, C^\ell) \to B(G, C)^\ell$$

in the following way. By the definition of $B(G, C)$, equation (3), we need to describe a map

$$\psi_1\colon L_{C^\ell} B(G^\ell) \times D^2 \to (L_C B(G) \times D^2)^\ell$$

that commutes with the evaluation maps. For this purpose, we let ψ be the product of two maps

$$\psi_2\colon L_{C^\ell} B(G^\ell) \xrightarrow{\cong} L_C B(G)^\ell, \quad \Delta\colon D^2 \to (D^2)^\ell.$$

The map ψ_2 is another natural isomorphism, while Δ is the diagonal embedding. With these choices, $\psi_1 = \psi_2 \times \Delta$ commutes with the evaluation map, completing the construction of ψ.

Lemma 4.5 *Let $\ell > 0$ be an integer and assume that C generates G and G is perfect. Then*:

(1) *C^ℓ generates G^ℓ.*

(2) *The Künneth theorem yields the isomorphism*

$$H_2(B(G, C)^\ell) \cong H_2(B(G, C))^\ell.$$

(3) *The map ψ commutes with the natural equivalence*

$$\widehat{R}_k(G^\ell, C^\ell) \cong \widehat{R}_k(G, C)^\ell.$$

Using (1), *this equivalence also commutes with the Schur invariant*:

$$\mathrm{sch}((f_1, f_2, \dots, f_\ell)) = (\mathrm{sch}(f_1), \mathrm{sch}(f_2), \dots, \mathrm{sch}(f_\ell)).$$

(4) *The induced map*

$$\psi_*\colon H_2(B(G^\ell, C^\ell)) \to H_2(B(G, C)^\ell) \cong H_2(B(G, C))^\ell$$

is injective.

Proof (1) C^ℓ generates a normal subgroup $N \le G^\ell$ that surjects onto each factor, so Lemma 2.3 tells us that $N = G^\ell$.

(2) Proposition 4.1(3) says that $B(G,C)$ is simply connected when C generates G, in particular that

$$H_1(B(G,C)) = 0.$$

Moreover, $H_0(B(G,C)) = \mathbb{Z}$ since $B(G,C)$ is connected. Thus the Künneth theorem simplifies to the stated isomorphism.

(3) The main step is to review the construction of a map

$$\phi\colon S^2 \to B(G,C)$$

representing $f \in \widehat{R}_k(G)$, and to then relate ϕ to the map ψ. Given

$$f\colon \pi_1(S^2 \smallsetminus [2k]) \to G,$$

we remove $2k$ open disks from S^2 instead of just $2k$ point punctures to obtain a surface $S^2 \smallsetminus k D^2$ with k boundary circles around the punctures. We can define a map

$$\phi\colon S^2 \smallsetminus [2k] \to B(G)$$

using the map f, and then use a fiber D^2 from the attachment $L_C B(G) \times D^2$ to extend ϕ across each puncture.

If we replace (G,C) by (G^ℓ, C^ℓ) in this construction, then it commutes with ψ because the same extension disk in S^2 is used ℓ times for each puncture. Moreover, if ψ_i is the i^{th} component of the map ψ, then the composition $\phi_i = \psi_i \circ \phi$ matches the i^{th} component f_i of f. Together with the Künneth isomorphism from the previous part of the lemma, this establishes that the i^{th} component of $\operatorname{sch}(f)$ is $\operatorname{sch}(f_i)$, as desired.

(4) Let $c \in C$ and consider the diagram

$$\begin{array}{ccccccc}
Z(c^{\times\ell})_{\text{ab}} & \xrightarrow{\kappa} & H_2(G^\ell) & \xrightarrow{\beta} & H_2(B(G^\ell, C^\ell)) & \xrightarrow{\sigma} & \mathbb{Z} \\
\downarrow\cong & & \downarrow\cong & & \downarrow\psi_* & & \downarrow\Delta \\
Z(c)^\ell_{\text{ab}} & \xrightarrow{\kappa^{\times\ell}} & H_2(G)^\ell & \xrightarrow{\beta^{\times\ell}} & H_2(B(G,C))^\ell & \xrightarrow{\sigma^{\times\ell}} & \mathbb{Z}^\ell
\end{array}$$

where each row is taken from Proposition 4.1(4) and is thus exact. Meanwhile the first vertical map is the elementary isomorphism from group theory; the second map is from the Künneth theorem and is an isomorphism because G is perfect; the third map is as indicated; and the fourth is the diagonal embedding of $\mathbb{Z}$ into $\mathbb{Z}^\ell$.

We claim that the diagram is commutative. Working from the left, the first square commutes by the definition of the map κ [38]. Given any group homomorphism $f: A \times B \to G$, there is always a biadditive map

$$f_{**}: A_{\mathrm{ab}} \times B_{\mathrm{ab}} \cong H_1(A) \times H_1(B) \to H_2(A \times B) \xrightarrow{f_*} H_2(G),$$

where the middle arrow is the Künneth map. The map κ is a linear restriction of f_{**}, where $A = Z(c)$ and $B = \langle c \rangle$. It is easy to confirm that κ for the group G^ℓ does the same thing as $\kappa^{\times \ell}$ for the group G. The second square is commutative by the way that ψ is constructed: since it is the identity on $B(G^\ell) = B(G)^\ell$, β and $\beta^{\times \ell}$ also do the same thing. Finally, the third square commutes because σ and $\sigma^{\times \ell}$ do the same thing by Proposition 4.1(5). By the Hurewicz theorem, we can represent any element of $H_2(B(G^\ell, C^\ell))$ by a map from S^2 and thus by a homomorphism

$$f: \pi_1(\Sigma_{a,b}) \to G^\ell.$$

Splitting this homomorphism f into ℓ homomorphisms to G, the difference $a - b$ is replicated ℓ times.

To complete the proof, since the diagram is commutative, the four lemma says that ψ_* is injective. □

4.4 An ultratransitivity theorem

Theorem 4.2 says that the action of $B_{k,k}$ on R_k^0 is transitive for all k large enough. Our goal now is Theorem 4.7, which, among other things, gives a complete description of this action when G is nonabelian simple. The structure of our argument is similar to one direction of the full monodromy theorem of Roberts and Venkatesh [37, Theorem 5.1]. However, Theorem 4.7 refines this special case of Roberts and Venkatesh in the same way that our prior result [28, Theorem 5.1] refines a result of Dunfield and Thurston [12, Theorem 7.4].

From here to the end of this article, we choose an element $c \in C$ and we declare the abbreviations

$$U := \mathrm{Aut}(G, c) \subseteq \widehat{U} := \mathrm{Aut}(G, C).$$

In this subsection we will only use $\widehat{U}$, but we will need the subgroup U soon enough. The group $\widehat{U}$ acts on $\widehat{R}_k^0$ because we can compose a homomorphism

$$f: \pi_1(D^2 \smallsetminus [2k]) \to G$$

with an element $\alpha \in \mathrm{Aut}(G, C)$. To see that $\widehat{U}$ preserves the property $\mathrm{sch}(f) = 0$, note that $\widehat{U}$ acts by group homomorphisms on $H_2(B(G, C))$ and $\mathrm{sch}(\alpha \circ f) = \alpha \cdot \mathrm{sch}(f)$. It acts freely on the subset R_k^0 because these homomorphisms are surjective. Moreover, the actions of $B_{k,k}$ and $\mathrm{Aut}(G, C)$ on R_k^0 commute. In other words, the corresponding permutation representation is a map

$$\rho\colon B_{k,k} \to \mathrm{Sym}_{\widehat{U}}(R_k^0).$$

Lemma 4.6 *Let G be a nonabelian simple group, let $C \subseteq G$ be a conjugacy class and let $\ell > 0$. Then $B_{k,k}$ eventually (as $k \to \infty$) acts $\widehat{U}$–set ℓ–transitively on R_k^0.*

Proof We choose k large enough that the conclusion of Theorem 4.2 holds for the finite group G^ℓ and the conjugacy class C^ℓ. Let

$$f_1, f_2, \dots, f_\ell \in R_k^0$$

lie in distinct $\widehat{U}$–orbits and consider the product homomorphism

$$f = f_1 \times f_2 \times \cdots \times f_\ell\colon \pi_1((D^2 \smallsetminus [2k])^\ell) \to G^\ell.$$

By Lemma 2.1, f is surjective. Since $\mathrm{sch}(f_j) = 0$ for all $j = 1, \dots, \ell$, Lemma 4.5 implies that $\mathrm{sch}(f) = 0$. If

$$e_1, e_2, \dots, e_\ell \in R_k^0$$

is another such list of homomorphisms with the same properties, with product e, then Theorem 4.2 says e and f are in the same orbit of $B_{k,k}$, as desired. □

Besides the map ρ already defined, let

$$\sigma_{J,E}\colon B_{k,k} \to \mathrm{Sym}(T_k(J, E))$$

be the action map of the braid group for every finite group J generated by a set of conjugacy classes $E \subseteq J$. Also let

$$\phi\colon B_{k,k} \to \mathrm{Sym}(k)^2$$

be the forgetful map that only remembers the permutation of the braid strands.

Theorem 4.7 *Let G be a finite, nonabelian, simple group and let $C \subseteq G$ be a conjugacy class. Then the image of $B_{k,k}$ under the joint homomorphism $\rho \times \sigma \times \phi$,*

$$\rho \times \phi \times \prod_{\substack{J \supseteq E \\ \#E < \#C}} \sigma_{J,E}\colon B_{k,k} \to \mathrm{Sym}_{\widehat{U}}(R_k^0) \times \mathrm{Sym}(k)^2 \times \prod_{\substack{J \supseteq E \\ \#E < \#C}} \mathrm{Sym}(T_k(J, E)),$$

eventually contains $\mathrm{Rub}_{\widehat{U}}(R_k^0)$, where here each group J is generated by a union of conjugacy classes $E \subseteq J$.

In other words, we can find a set of *pure* braids that act by the entire Rubik group on R_k^0, while simultaneously acting trivially on every $T_k(J,E)$ with $\#E < \#C$. It follows that such braids also act trivially on the set $\widehat{R}_k \smallsetminus R_k$ of nonsurjective maps in $\widehat{R}_k$.

Proof Following Dunfield and Thurston [12] and Roberts and Venkatesh [37], we use the corollary of the classification of finite simple groups that every 6–transitive permutation group on a finite set is ultratransitive. Lemma 4.6 shows $B_{k,k}$ eventually acts $\mathrm{Aut}(G,C)$–set 6–transitively on R_k^0. It follows that $B_{k,k}$ eventually acts 6–transitively (in the usual sense) on $R_k^0/\mathrm{Aut}(G,C)$. By Theorem 2.5, the image of ρ contains $\mathrm{Rub}_{\widehat{U}}(R_k^0)$.

For the rest of the properties of the joint homomorphism, Lemma 2.2 tells us it is enough to show that $\mathrm{Rub}_{\widehat{U}}(R_k^0)$ does not have any simple quotients that are subquotients of

$$\mathrm{Sym}(k)^2 \times \prod_{\substack{J \supseteq E \\ \#E < \#C}} \mathrm{Sym}(T_k(J,E)).$$

By Lemmas 2.6 and 2.4, it suffices to show that $\mathrm{Alt}(R_k^0/\widehat{U})$ is not a subquotient of $\mathrm{Sym}(T_k(J,E))$ for any group J generated by a union of conjugacy classes E, nor a subquotient of $\mathrm{Sym}(k)$. To prove this, we show that $\mathrm{Alt}(R_k^0/\widehat{U})$ is eventually larger than any of these other groups. This follows from comparing $\#T_k(J,E) = \#E^{2k}$ to the bound in Lemma 4.4, which shows that

$$\lim_{k\to\infty} (\#R_k^0)^{1/2k} = \#C.$$

Meanwhile $\mathrm{Sym}(k)$ by definition acts on a set that only grows linearly in k, not exponentially. □

Remark Although our proof of Theorem 4.7 (hence also our main theorem) depends on the classification of finite simple groups via the 6–transitivity corollary, we conjecture that the classification can be avoided. The analogous step in our previous work [28] is a result of Dunfield and Thurston [12, Theorem 7.4], which they argue in the same way. However, they point out that they could use a nonclassification result of Dixon and Mortimer [11, Theorem 5.5B], which says that a permutation group on a finite set is ultratransitive if it is both 2–transitive, and locally ℓ–transitive on a single subset of size $\ell = \Omega(\log k)$. They show that this transitivity theorem suffices (for mapping class groups of unmarked, closed surfaces) when the Schur invariant vanishes thanks to a result of Gilman [17], but the argument can be extended to any value of the Schur invariant. It would suffice to find an analogue of Gilman's theorem for braid groups.

5 Proof of Theorem 1.1

As before, let G be a nonabelian simple group, let $C \subseteq G$ be a nontrivial conjugacy class (which necessarily generates G) and let $c \in C$ be a distinguished element. Again, let

$$U := \mathrm{Aut}(G, c) \subseteq \mathrm{Aut}(G, C) := \widehat{U}.$$

Also for this section, choose some fixed k large enough for the conclusions of Theorem 4.7.

Remark Although our proof of Theorem 1.1 is similar to that of our previous result for mapping class groups [28], we will face a new technical difficulty. Namely, even though the available braid actions are $\widehat{U}$–equivariant, the reduction from ZSAT is locally only U–equivariant. We will define a zombie symbol z using the distinguished element c, which is not a $\widehat{U}$–invariant choice. If we represented z using all of C or in any other $\widehat{U}$–invariant manner, then the construction could only produce an intractable link invariant rather than specifically a knot invariant.

5.1 Alphabets and gadgets

In this subsection, we will define a U–set alphabet A with subsets $I, F \subseteq A$ and a zombie symbol $z \in I \cap F$ such that $\mathsf{ZSAT}_{U,A,I,F}$ satisfies Theorem 3.1 and is thus almost parsimoniously #P–complete, for use to the end of this article. Then we will define pure braid gadgets, which we will later use to replace gates in a ZSAT circuit. Here a *gadget* is a semirigorous concept in theoretical computer science, by definition a local combinatorial replacement to implement a complexity reduction. Our gadget to replace one gate will be a braid with a fixed number of strands. We will later concatenate these braid gadgets to replace an entire circuit with a braid with a linear number of strands.

Let the zombie symbol be

$$z := (c, c, \dots, c) \in \widehat{R}_k^0,$$

and let the alphabet be

$$A := \{z\} \cup \{(g_1, g_2, \dots, g_{2k}) \in R_k^0 \mid g_1 = g_{2k} = c\}.$$

That is, the nonzombie symbols in A are surjections with trivial Schur invariant such that the first and last punctures map to c specifically. The initialization and finalization

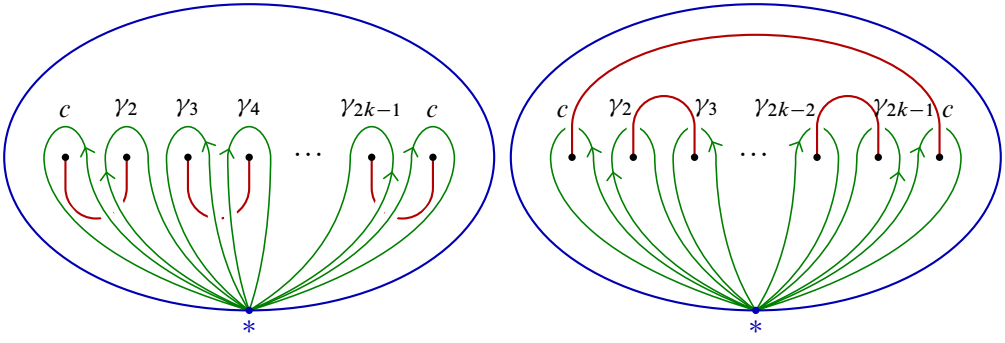

Figure 3: These two sets of red plats impose the initial and final constraints, respectively.

conditions are specified by restricting to homomorphisms that factor through the two trivial tangles in Figure 3, respectively. Precisely, we define the initial and final subalphabets by

$$I := \{(g_1, g_2, \dots, g_{2k}) \in A \mid g_{2i} = g_{2i-1} \text{ for all } i \leq k\},$$
$$F := \{(g_1, g_2, \dots, g_{2k}) \in A \mid g_{2i} = g_{2i+1} \text{ for all } i \leq k-1\}.$$

It is straightforward to verify that U, A, I and F satisfy the conditions of Theorem 3.1.

We now construct braid gadgets that simulate gates in $\mathrm{Rub}_U(A^2)$. Let $D^2 \smallsetminus [4k]$ be a pointed disk with $4k$ punctures, and divide it into two half-disks by a straight line, so that each half contains the basepoint and half of the $4k$ punctures, as in Figure 4. This allows us to identify $T_k \times T_k \cong C^{2k} \times C^{2k}$ with $T_{2k} \cong C^{4k}$. It is straightforward to verify that this identification takes $R_k^0 \times R_k^0$ to a subset of R_{2k}^0. In particular, we identify A^2 with a subset of $R_{2k}^0 \cup \{(z, z)\}$.

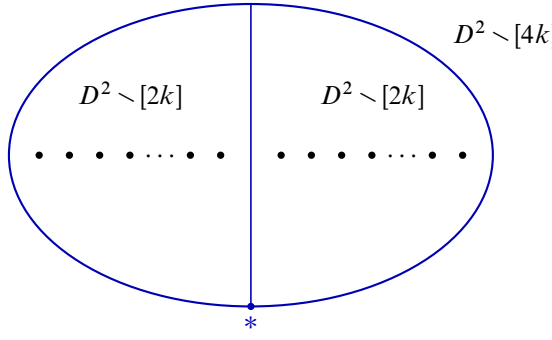

Figure 4: Splitting $D^2 \smallsetminus [4k]$ as a boundary sum of two copies of $D^2 \smallsetminus [2k]$.

Corollary 5.1 *For every gate $\delta \in \mathrm{Rub}_U(A^2)$, there is a braid word $b(\delta)$, interpreted also as a braid element $b(\delta) \in B_{2k,2k}$, with the following properties*:

(1) *$b(\delta)$ acts on A^2 as δ.*

(2) *$b(\delta)$ acts trivially on $R^0_{2k} \smallsetminus \widehat{U} \cdot (A^2)$.*

(3) *$b(\delta)$ acts trivially on $T_{2k}(J, E) \cong E^{4k}$ for every group J generated by a union of conjugacy classes E with $\#E < \#C$.*

(4) *$b(\delta)$ is a pure braid, ie $b(\delta) \in PB_{4k} \leq B_{2k,2k}$.*

The existence of $b(\delta)$ follows immediately from Theorem 4.7. In fact, for each δ in $\mathrm{Rub}_U(A^2)$ there are infinitely many braids that satisfy properties (1)–(4), but it is important for our reduction that we fix some suitable $b(\delta)$ for each δ.

Note that property (1) specifies the action of $b(\delta)$ on A^2, but it implies more than that, because the action of $B_{2k,2k}$ is $\widehat{U}$–equivariant while A^2 is only closed under the action of U. This action has a unique $\widehat{U}$–equivariant extension to $\widehat{U} \cdot (A^2)$. Meanwhile, property (3) implies that $b(\delta)$ acts trivially on $\widehat{R}_{2k} \smallsetminus R_{2k}$, so together the first three properties specify all of the action of $b(\delta)$ on $\widehat{R}^0_{2k}$. However, $b(\delta)$ is not fully specified on all of $\widehat{R}_{2k}$, because we place no restrictions on its effect on $f \in R^s_{2k}$ with nonvanishing Schur invariant, $s \neq 0$.

We record as a lemma several invariance properties of $b(\delta)$ that we have already discussed, either here or previously.

Lemma 5.2 *If $\delta \in \mathrm{Rub}_U(A^2)$, then $b(\delta) \in PB_{4k}$ acting on $\widehat{R}_{2k}$ preserves all of the sets*

$$\{(z,z)\} \subsetneq A^2 \subsetneq \widehat{U} \cdot (A^2) \subsetneq (\widehat{U} \cdot A)^2 \subsetneq (\widehat{R}^0_k)^2 \subsetneq \widehat{R}^0_{2k} \subsetneq \widehat{R}_{2k}.$$

Note that it is easy to confuse the set A^2 with the slightly larger $\widehat{U} \cdot (A^2)$ and $(\widehat{U} \cdot A)^2$, and the set $(\widehat{R}^0_k)^2$ with the slightly larger $\widehat{R}^0_{2k}$. In the proof, it will be crucial that each $b(\delta)$ preserves both A^2 and $(\widehat{R}^0_k)^2$.

5.2 The reduction

Let Z be an instance of $\mathsf{ZSAT}_{U,A,I,F}$, with U, A, I and F as in Section 5.1. Recall this means Z is a planar U–equivariant reversible circuit over the alphabet A. Suppose that Z has width n, so that it acts on n symbols; and length ℓ, so that it has ℓ gates.

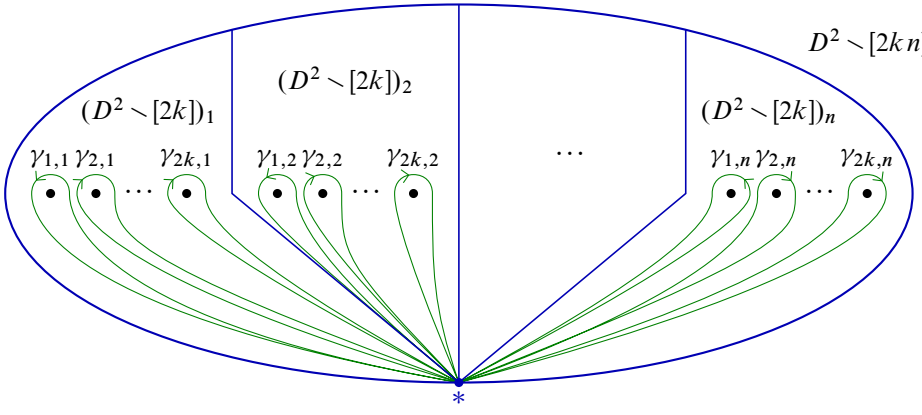

Figure 5: Punctured disks that encode n symbols of a ZSAT circuit.

Consider the disk $D^2 \smallsetminus [2kn]$ with $2kn$ punctures and a separate basepoint $* \in \partial D^2$. Divide it into n disks $(D^2 \smallsetminus [2k])_i$ with $1 \le i \le n$ so that each one contains the basepoint, as indicated in Figure 5. Also pick generators $\{\gamma_{j,i}\}$ for each $\pi_1((D^2 \smallsetminus [2k])_i)$ as indicated in the figure, where $1 \le j \le 2k$.

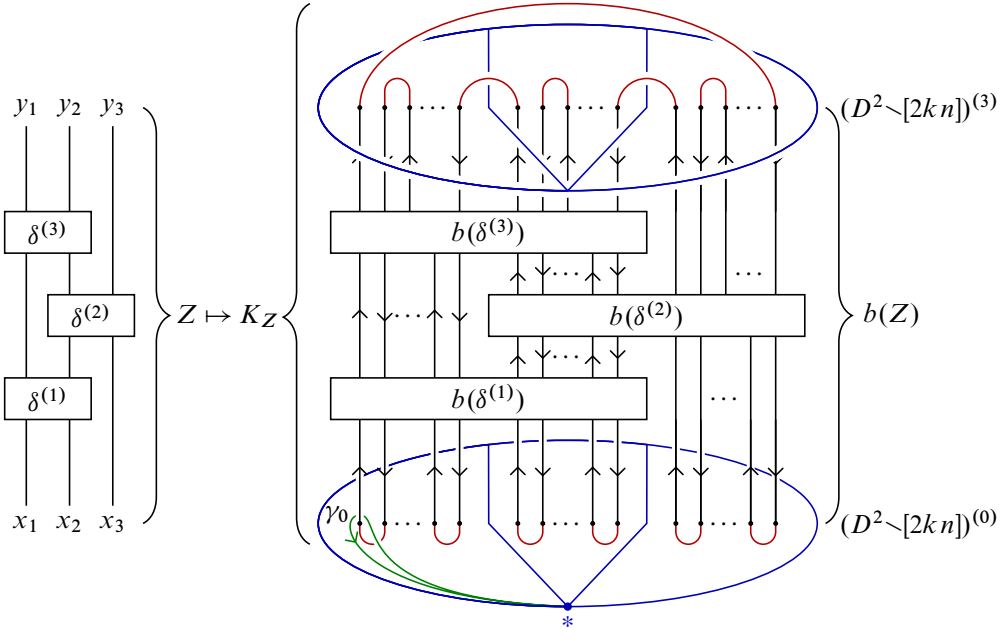

Figure 6: Reducing a circuit Z with $n = 3$ variables to the knot K_Z.

As in Figure 6, we convert Z into a braid diagram $b(Z)$ by replacing each strand in Z with $2k$ parallel strands and each gate $\delta^{(m)}$ in Z with the braid gadget $b(\delta^{(m)})$, where here $1 \le m \le \ell$. Let $K(Z)$ be the oriented link diagram formed by the plat closure of $b(Z)$ indicated in the figure, and, for each m with $0 \le m \le \ell$, let $(D^2 \smallsetminus [2kn])^{(m)}$ be a disk transverse to the braid, so that these disks and the braid gadgets alternate. Each disk $(D^2 \smallsetminus [2kn])^{(m)}$ is also divided into subdisks $\{(D^2 \smallsetminus [2kn])_i^{(m)}\}$ as before, with loops $\{\gamma_{j,i}^{(m)}\}$. Finally, let $\gamma_0 \in \pi_1(S^3 \smallsetminus K(Z))$ be the indicated meridian. In other words, $\gamma_0 = \gamma_{1,1}^{(0)}$.

We are interested in homomorphisms

$$f : \pi_1(S^3 \smallsetminus K(Z)) \to G$$

such that $f(\gamma_0) = c$. Using the system of disks and loops just defined, we can restrict f to other maps and elements as follows:

$$\begin{aligned} f^{(m)} &: \pi_1((D^2 \smallsetminus [2kn])^{(m)}) \to G, \\ f_i^{(m)} &: \pi_1((D^2 \smallsetminus [2k])_i^{(m)}) \to G, \\ f_{j,i}^{(m)} &= f(\gamma_{j,i}^{(m)}) \in C. \end{aligned}$$

We can also write $f_i^{(m)} \in T_k \cong C^{2k}$, and we can think of the map $f_i^{(m)}$ as a list of the group elements $(f_{j,i}^{(m)})_{j=1}^{k}$. For simplicity we rename the first and last levels of f:

$$p := f^{(0)}, \quad q := f^{(\ell)}.$$

The inclusion map

$$\imath_* : \pi_1((D^2 \smallsetminus [2kn])^{(0)}) \to \pi_1(S^3 \smallsetminus K(Z))$$

is always a surjection and never a bijection. Our goal is to show that a map p from the former extends to a map f from the latter if and only if p corresponds to a solution to the circuit Z with $q = Z(p)$. (Moreover, we show that there are no nontrivial solutions if we replace G with a group J generated by a smaller conjugacy class.)

Lemma 5.3 *Let Z be an instance of $\mathsf{ZSAT}_{U,A,I,F}$ and let $\#Z$ denote the number of solutions to Z. Then the diagram $K(Z)$ and meridian γ_0 have the following properties*:

(1) *$K(Z)$ is a knot.*

(2) *If J is a noncyclic group generated by a conjugacy class E with $\#E < \#C$, then $\#Q(K(Z); J, E) = 0$.*

(3) $\#H(K(Z), \gamma_0; G, c) = \#Z$.

Proof (1) Every braid gadget $b(\delta)$ as in Corollary 5.1 is a pure braid, so our choice of plats in Figure 6 guarantees that $K(Z)$ is a knot, rather than a link with more than one component.

(2) Let J be a group generated by a conjugacy class E such that $\#E < \#C$, and retain the notation f, p and q defined above for the group G. Following Corollary 5.1, an arbitrary gadget $b(\delta)$ acts on $T_{2k}(J, E)$ by definition, and acts trivially by construction. In particular, each braid gadget $b(\delta^{(m)})$ acts on some pair $(f_i^{(m-1)}, f_{i+1}^{(m-1)})$, and does nothing to that pair. Thus, for the purpose of computing either $\#H(K(Z); J, E)$ or $\#Q(K(Z); J, E)$, $K(Z)$ is equivalent to the unknot. Since by hypothesis J is not cyclic, we obtain $\#Q(K(Z); J, E) = 0$, as desired.

(3) Let $X(Z)$ be the set of solutions to the circuit Z. We will show that $X(Z) = H(K(Z), \gamma_0; G, c)$ in the natural sense. If $q = Z(p)$ is a solution to Z, then, by definition,

$$(p_1, p_2, \dots, p_n) \in I^n, \quad Z(p) = (q_1, q_2, \dots, q_n) \in F^n.$$

By Corollary 5.1, each braid gadget $b(\delta)$ acts on A^2 exactly as δ does, and therefore the braid $b(Z)$ acts on A^n exactly as the circuit Z does. By the definition of the initial subalphabet I, the map $p = f^{(0)}$ factors through the plat attached to the bottom of the braid $b(Z)$. Meanwhile, the definition of the alphabet A together with the definition of the final subalphabet F together imply that $q = b(Z) \cdot p$ factors through the plat attached to the top of $b(Z)$. Most of the U-turns at the top of the plat are internal to one symbol $q_i \in A$, and these force $q_i \in F$. The others connect either $q_{2k,i}$ with $q_{1,i+1}$ or $q_{2k,n}$ with $q_{1,1}$. These constraints hold automatically in the alphabet A, because they reduce to the equation $c = c$. Finally, $p_1 \in A$ also gives us that $f(\gamma_0) = c$. This establishes that $X(Z) \subseteq H(K(Z), \gamma_0; G, c)$. In fact it establishes a little more, namely that any other element of $H(K(Z), \gamma_0; G, c)$ cannot come from $p = f^{(0)} \in A^n$.

Conversely, let $f \in H(K(Z), \gamma_0; G, c) \smallsetminus X(Z)$ be a hypothetical spurious homomorphism. Then, tautologically, $p = f^{(0)} \in T_k^n \cong (C^{2k})^n$, but we quickly obtain an important restriction. Each p_i factors through the initial plat attached to $(D^2 \smallsetminus [2k])_i^{(0)}$, so Lemma 4.3 tells us that $\mathrm{sch}(p_i) = 0$ and thus that $p_i \in \widehat{R}_k^0$. Moreover, Lemma 5.2 tells us that every braid gadget preserves this condition, so $f_i^{(m)} \in \widehat{R}_k^0$ for every m and i. In other words, we can interpret $b(Z)$ as a circuit that uses the larger alphabet $\widehat{R}_k^0 \supseteq A$.

We claim that we can further restrict the alphabet to $\widehat{U} \cdot A$. If $p_i = f_i^{(0)} \in \mathbb{R}_k^0 \smallsetminus \widehat{U} \cdot A$ for some i, then Corollary 5.1 also tells us that no gate gadget changes this value, so

that, in particular, $q_i = p_i$. But then the initial and final plat closures together tell us that

$$p_i = (e, e, \ldots, e)$$

for some $e \in C$, which thus means that

$$p_i = q_i \in \mathrm{Inn}(G) \cdot \{z\} \subseteq \widehat{U} \cdot A$$

after all. By Lemma 5.2, the condition that $p = f^{(0)} \in (\widehat{U} \cdot A)^n$ is also preserved through every gate gadget in $b(Z)$.

We now show that $f^{(m)} \in \widehat{U} \cdot A^n$ for every m. The condition that $f_i^{(m)} \in \widehat{U} \cdot A$ tells us that each symbol $f_i^{(m)}$ begins and ends with the same group element $e \in C$, and what we would like to know is that e does not depend on i. The final plat closure makes this immediate for $q = f^{(\ell)}$, and then Lemma 5.2 tells us that the condition is preserved in reverse $f^{(m)}$ as m decreases.

Finally, because $p = f^{(0)} \in \widehat{U} \cdot A^n$ and $f(\gamma_0) = c$, we conclude that $p \in A^n$. □

To conclude the proof of Theorem 1.1, the knot $K(Z)$ can be constructed from Z in polynomial time as a function of the number of gates in Z, since it is just direct replacement of each gate δ by the corresponding gate gadget $b(\delta)$. Although we do not know how large k must be for Corollary 5.1 to hold, for any fixed G and C it is a constant amount of effort to find such a k and then compute a fixed $b(\delta)$ for each δ. Thus it is a parsimonious reduction from $\#\mathsf{ZSAT}_{U,A,I,F}$ to $\#H(-, G, c)$ that preserves the #P–completeness properties stated in Theorem 3.1.

Remark As in our previous work [28], the proof of Theorem 1.1 establishes an efficient bijection between $Q(K(Z), \gamma_0; G, c)$ and the orbits of nontrivial solutions to $\#\mathsf{ZSAT}_{U,A,I,F}$, and therefore the set of certificates in any problem in #P. This conclusion is a refinement of parsimonious reduction which is known as *Levin reduction*.

6 Open problems

As with our previous theorem about homology 3–spheres [28], we conjecture that $\#Q(K; G, C)$ is also computationally intractable when K is a randomly chosen knot. There are various inequivalent models for choosing a knot at random [14], and we believe that $\#Q(K; G, C)$ should be intractable for many of them. Hardness in random cases is a known property for some #P–complete problems [9].

Also in our previous work, we first had in mind that the analogous invariant $\#Q(M;G)$ is intractable for 3–manifolds M; later we sharpened the construction to make M a homology 3–sphere. Theorem 1.1 is in keeping with the analogy that a homology 3–sphere is like a knot, while a general 3–manifold is like a link. However, a deeper analogy is that a homology 3–sphere, among 3–manifolds, is like a knot with trivial Alexander polynomial, among knots. We conjecture that Theorem 1.1 also holds for knots with trivial Alexander polynomial. This would better motivate the restriction that G should be a nonabelian simple group.

References

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

[2] **I Agol**, **J Hass**, **W Thurston**, *The computational complexity of knot genus and spanning area*, Trans. Amer. Math. Soc. 358 (2006) 3821–3850 MR Zbl

[3] **S Arora**, **B Barak**, *Computational complexity: a modern approach*, Cambridge Univ. Press (2009) MR Zbl

[4] **R Barbanchon**, *On unique graph 3–colorability and parsimonious reductions in the plane*, Theoret. Comput. Sci. 319 (2004) 455–482 MR Zbl

[5] **C H Bennett**, **E Bernstein**, **G Brassard**, **U Vazirani**, *Strengths and weaknesses of quantum computing*, SIAM J. Comput. 26 (1997) 1510–1523 MR Zbl

[6] **H L Bodlaender**, *A partial k–arboretum of graphs with bounded treewidth*, Theoret. Comput. Sci. 209 (1998) 1–45 MR Zbl

[7] **M Bordewich**, **M Freedman**, **L Lovász**, **D Welsh**, *Approximate counting and quantum computation*, Combin. Probab. Comput. 14 (2005) 737–754 MR Zbl

[8] **N Brand**, *Classifying spaces for branched coverings*, Indiana Univ. Math. J. 29 (1980) 229–248 MR Zbl

[9] **J-Y Cai**, **A Pavan**, **D Sivakumar**, *On the hardness of permanent*, from "STACS 99" (C Meinel, S Tison, editors), Lect. Notes Comput. Sci. 1563, Springer (1999) 90–99 MR Zbl

[10] **R H Crowell**, **R H Fox**, *Introduction to knot theory*, Ginn, Boston (1963) MR Zbl

[11] **J D Dixon**, **B Mortimer**, *Permutation groups*, Grad. Texts Math. 163, Springer (1996) MR Zbl

[12] **N M Dunfield**, **W P Thurston**, *Finite covers of random 3–manifolds*, Invent. Math. 166 (2006) 457–521 MR Zbl

[13] **J S Ellenberg**, **A Venkatesh**, **C Westerland**, *Homological stability for Hurwitz spaces and the Cohen–Lenstra conjecture over function fields, II*, preprint (2012) arXiv:1212.0923v1

[14] **C Even-Zohar**, *Models of random knots*, J. Appl. Comput. Topol. 1 (2017) 263–296 MR Zbl

[15] **R H Fox**, *A quick trip through knot theory*, from "Topology of 3–manifolds and related topics" (M K Fort, editor), Prentice-Hall, Englewood Cliffs, NJ (1962) 120–167 MR Zbl

[16] **M D Fried**, **H Völklein**, *The inverse Galois problem and rational points on moduli spaces*, Math. Ann. 290 (1991) 771–800 MR Zbl

[17] **R Gilman**, *Finite quotients of the automorphism group of a free group*, Canadian J. Math. 29 (1977) 541–551 MR Zbl

[18] **E Goursat**, *Sur les substitutions orthogonales et les divisions régulières de l'espace*, Ann. Sci. École Norm. Sup. 6 (1889) 9–102 MR Zbl

[19] **P Hall**, *The Eulerian functions of a group*, Q. J. Math. 7 (1936) 134–151 Zbl

[20] **P de la Harpe**, **V F R Jones**, *Graph invariants related to statistical mechanical models: examples and problems*, J. Combin. Theory Ser. B 57 (1993) 207–227 MR Zbl

[21] **R Impagliazzo**, **R Paturi**, *On the complexity of k–SAT*, J. Comput. System Sci. 62 (2001) 367–375 MR Zbl

[22] **F Jaeger**, *Tutte polynomials and link polynomials*, Proc. Amer. Math. Soc. 103 (1988) 647–654 MR Zbl

[23] **F Jaeger**, **D L Vertigan**, **D J A Welsh**, *On the computational complexity of the Jones and Tutte polynomials*, Math. Proc. Cambridge Philos. Soc. 108 (1990) 35–53 MR Zbl

[24] **M Jerrum**, **A Sinclair**, **E Vigoda**, *A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries*, J. ACM 51 (2004) 671–697 MR Zbl

[25] **D Koenig**, **A Tsvietkova**, *NP-hard problems naturally arising in knot theory*, preprint (2018) arXiv:1809.10334

[26] **H Krovi**, **A Russell**, *Quantum Fourier transforms and the complexity of link invariants for quantum doubles of finite groups*, Comm. Math. Phys. 334 (2015) 743–777 MR Zbl

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

[28] **G Kuperberg**, **E Samperton**, *Computational complexity and 3–manifolds and zombies*, Geom. Topol. 22 (2018) 3623–3670 MR Zbl

[29] **M Lackenby**, *Some conditionally hard problems on links and 3–manifolds*, Discrete Comput. Geom. 58 (2017) 580–595 MR Zbl

[30] **R J Lipton**, **R E Tarjan**, *A separator theorem for planar graphs*, SIAM J. Appl. Math. 36 (1979) 177–189 MR Zbl

[31] **D Lokshtanov**, **D Marx**, **S Saurabh**, *Lower bounds based on the exponential time hypothesis*, Bull. Eur. Assoc. Theor. Comput. Sci. 105 (2011) 41–71 MR Zbl

[32] **A de Mesmay**, **Y Rieck**, **E Sedgwick**, **M Tancer**, *The unbearable hardness of unknotting*, from "35th International Symposium on Computational Geometry" (G Barequet, Y Wang, editors), Leibniz Int. Proc. Inform. 129, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern (2019) art. id. 49 MR Zbl

[33] **J Munkres**, *Algorithms for the assignment and transportation problems*, J. Soc. Indust. Appl. Math. 5 (1957) 32–38 MR Zbl

[34] **K Reidemeister**, *Knoten und Gruppen*, Abh. Math. Sem. Univ. Hamburg 5 (1927) 7–23 MR Zbl

[35] **K A Ribet**, *On l–adic representations attached to modular forms*, Invent. Math. 28 (1975) 245–275 MR Zbl

[36] **K A Ribet**, *Galois action on division points of abelian varieties with real multiplications*, Amer. J. Math. 98 (1976) 751–804 MR Zbl

[37] **D P Roberts**, **A Venkatesh**, *Hurwitz monodromy and full number fields*, Algebra Number Theory 9 (2015) 511–545 MR Zbl

[38] **E Samperton**, *Schur-type invariants of branched G–covers of surfaces*, from "Topological phases of matter and quantum computation" (P Bruillard, C Ortiz Marrero, J Plavnik, editors), Contemp. Math. 747, Amer. Math. Soc., Providence, RI (2020) 173–197 MR Zbl

[39] **L G Valiant**, *The complexity of computing the permanent*, Theoret. Comput. Sci. 8 (1979) 189–201 MR Zbl

[40] **L G Valiant**, **V V Vazirani**, *NP is as easy as detecting unique solutions*, Theoret. Comput. Sci. 47 (1986) 85–93 MR Zbl

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

Department of Mathematics, University of Illinois at Urbana-Champaign
Urbana, IL, United States

greg@math.ucdavis.edu, smprtn@illinois.edu

Received: 20 November 2019 Revised: 6 September 2020

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2021 is US $560/year for the electronic version, and $835/year (+$65, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 7 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Berkeley, CA 94704, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

AGT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 21 Issue 3 (pages 1075–1593) 2021