Title

Coloring invariants of knots and links are often intractable

Journal

Algebraic & Geometric Topology, 21(3)

ISSN

1472-2739

Authors

Kuperberg, Greg

Samperton, Eric

Publication Date

2021-08-11

DOI

10.2140/agt.2021.21.1479

Coloring invariants of knots and links are often intractable

Greg Kuperberg*
University of California, Davis

Eric Samperton†
University of California, Santa Barbara

(Dated: July 16, 2019)

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.

1. INTRODUCTION

Let K be an oriented knot in the 3-sphere S^3 described by some given knot diagram. Fox [7, Exer. VI.6-7] popularized the idea of a 3-coloring of the diagram K, which is now also called a Fox coloring [8]. 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 $\pi_1(S^3 \smallsetminus K)$ [23]: 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

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

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.

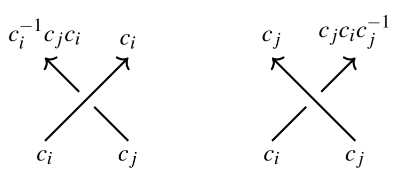

Figure 1. The Wirtinger relations.

If $G = S_3$ is the symmetric group on 3-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" [14]. Our main result, Theorem 1.1, demonstrates that Fox's frustration was prescient; see Section 1.1.

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

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

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 non-abelian simple, then $\#Q(K;G,C) = \#Q(K,\gamma;G,c)$ is computationally intractable, and remains so even when every homomorphism $f \in H(K;G,C)$ is promised either to be surjective or have cyclic image.

Theorem 1.1. *Let G be a fixed, finite, non-abelian 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 #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.

* greg@math.ucdavis.edu; Partly supported by NSF grant CCF-1716990
† eric@math.ucsb.edu; Partly supported by NSF grant CCF-1716990

Before reviewing the definition of #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

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

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

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

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

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

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

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

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 [21, Sec. 2.1], as well as Arora-Barak [2] and the Complexity Zoo [30]. 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 $c : A^* \to \mathbb{N}$ given by the equation

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

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

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

is a predicate that can be computed in polynomial time. A counting problem $c \in$ #P is *parsimoniously #P-complete* when every problem $b \in$ #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 $f : A^* \to A^*$ 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 that 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 [29]. 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 $m(\Gamma) = 0$ can be computed in polynomial time [22], and $m(\Gamma)$ can be approximated in randomized polynomial time [19]. 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 [4].

The analogous concepts for existence questions are the complexity class NP and the NP-completeness property. A decision function $d : A^* \to \{\text{yes}, \text{no}\}$ 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 $e \in$ NP, $e(x) = d(f(x))$ for some f computable in polynomial time.

In particular, Theorem 1.1 implies NP-completeness results for the existence of C-colorings. De Mesmay, Rieck, Sedgwick, and Tancer [9] cite us for the first known NP-hardness result for knots in S^3 (as opposed to knots in general 3-manifolds [1]), so we record this as an explicit corollary.

Corollary 1.2. *Let G be a fixed, finite, non-abelian 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 deciding whether $Q(K;G,C)$ is non-empty or $\#H(K,G,C) > \#C$ is* NP*-complete via Post-Karp reduction.*

Note also that much more is true thanks to a result of Valiant and Vazirani [28]: Distinguishing *any* two values of a parsimoniously #P-hard problem is NP-hard with randomized reduction [21, Thm. 2.1].

Some partial information about the unadjusted counting invariant $\#H(K;G,C)$ can be computed efficiently; for instance, that it always at least $\#C$. However, Theorem 1.1 and equation (1) together imply that this extra information can be trivial. We call a counting problem $c \in$ #P *almost parsimoniously* #P*-complete* if for every $b \in$ #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 #P-complete [3]. Likewise, Theorem 1.1 shows $\#H(K;G,C)$ is almost parsimoniously #P-complete.

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 #P-complete for any fixed $m \geq 5$ and any fixed conjugacy class C of permutations with at least 4 fixed points [20]. 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 $\#Q(L;A_m,C) > 0$.

1.2. Outline of the proof

Our proof of Theorem 1.1 follows our proof of the analogous theorem for homology 3-spheres [21], 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 $\#\mathsf{P}$ by the same argument as in the 3-manifold case [21, Thm. 2.7].

The reduction begins with a counting version of circuit satisfiability, #CSAT, that is rather directly parsimoniously #P-complete [21, Thm. 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 $D^2 \setminus [n]$ 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

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

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 non-surjective 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 [15] that shows that the action is at least transitive when G is any finite group. This refinement first appeared in version 1 of a retracted e-print of Ellenberg, Venkatesh, and Westerland [12, Thm. 7.6.1]; the second author of this article later found a topological proof [27, Thm. 1.1]. The second ingredient is a set of surjectivity results for group homomorphisms (Section 2).

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 [17] and Ribet's Lemma [24, 25]. In our research, we first saw it stated by Dunfield and Thurston [11, Lem. 3.7].

Lemma 2.1 (After Goursat-Ribet [25, Lems. 5.2.1 & 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-Thurston assume that the target groups are finite even though their proofs do not use this hypothesis. Dunfield-Thurston also state that the result is due to Hall [18]. 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 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 [21, Sec. 3.2]. The following is also an adaptation of Goursat's Lemma.

Lemma 2.2 (After Goursat [21, Lem. 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 $f(B)$ contains G_1.

Finally we will need the following two related lemmas.

Lemma 2.3 (Ribet [25, Sec. 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 ([21, Lem. 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

Let G be a group and let X be a G-set with finitely many orbits. 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

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

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

$$\mathrm{Sym}(n,G) \stackrel{\text{def}}{=} G \,\mathrm{wr}_m\, \mathrm{Sym}(n) = G^{\times n} \rtimes \mathrm{Sym}(n).$$

Likewise let

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

We need two results about Rubik groups from our previous work [21] 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*, by definition 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 ([21, Thm. 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 ([21, Lem. 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 [21].

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 non-trivial 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 $\mathrm{Rub}_G(A^2)$. (Remark: 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

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

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 ([21, Lem. 4.1])**.** *$\#\mathsf{ZSAT}_{G,A,I,F}$ 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 $(z,\ldots,z)$. 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.

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 [26, Thm. 5.1].

4.1. Conjugacy-restricted homomorphisms and colored braid subgroups

Recall that $D^2 \setminus [2k]$ 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 base point $*$ that is not on the same line. Also as shown, choose a list of generators of $\pi_1(D^2 \setminus [2k])$ represented by simple closed curves $\gamma_1, \ldots, \gamma_{2k}$, where each γ_i winds counterclockwise around the puncture p_i when it is positive (when i is odd) and clockwise

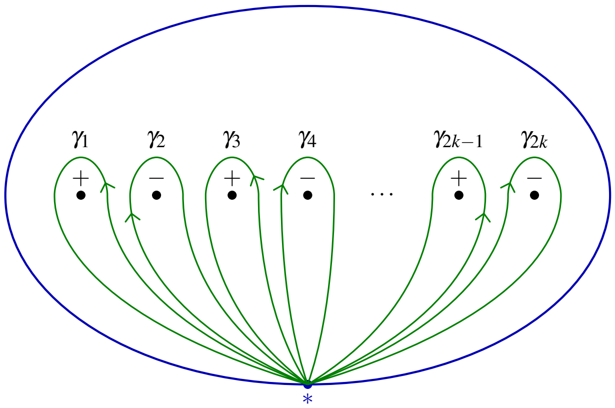

Figure 2. Generators for $\pi_1(D^2 \setminus [2k])$.

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, *e.g.*, $\gamma_1\gamma_2$ is the element of $\pi_1(D^2 \setminus [2k])$ 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 $\pi_1(D^2 \setminus [2k])$ to G:

$$T_k(G,C) \stackrel{\text{def}}{=} \{f : \pi_1(D^2 \setminus [2k]) \to G \mid f(\gamma_i) \in C\}$$
$$\hat{R}_k(G,C) \stackrel{\text{def}}{=} \{f \in T_k(G,C) \mid f(\gamma_\infty) = 1\}$$
$$R_k(G,C) \stackrel{\text{def}}{=} \{f \in \hat{R}_k(G,C) \mid f \text{ is onto}\}.$$

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

$$T_k \stackrel{\text{def}}{=} T_k(G,C) \qquad \hat{R}_k \stackrel{\text{def}}{=} \hat{R}_k(G,C) \qquad R_k \stackrel{\text{def}}{=} R_k(G,C).$$

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

Since $\pi_1(D^2 \setminus [2k])$ is freely generated by $\gamma_1, \gamma_2, \ldots, \gamma_{2k}$, 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 $\hat{R}_k$ and R_k, each such f factors through $\pi_1(S^2 \setminus [2k])$, where $S^2 \setminus [2k]$ is a punctured sphere obtained by collapsing the boundary ∂D^2 to a point. Also by abuse of notation, we will also interpret each such f as having this domain instead.

The homomorphism sets $T_k(G,C)$, $\hat{R}_k(G,C)$, and $R_k(G,C)$ are all invariant under the colored braid group $B_{k,k} \le B_{2k}$, 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, *i.e.*, if it is true for all k large enough.

The main tool that we need is the Brand classifying space $B(G,C)$ [5]. 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 [27]. Before stating these properties, we first review the definition. The free loop space $LB(G)$ comes with an evaluation map

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

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:

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

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 base point 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 [12], Roberts and Venkatesh [26] extend the notation $H_*(G)$ in an ad hoc way to a reduced Schur multiplier that they denote "$H_2(G,C)$". 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.

Proposition 4.1 ([27])**.** *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 $\Sigma_{a,b} = S^2 \setminus [a+b]$ be a punctured sphere with a points marked $+$, b points marked $-$, and a base point $*$. 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 part 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*

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

where L is a tangle.

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

4. *There is an exact sequence*

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

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 $H_2(B(G,C))_{\mathrm{tor}}$.

5. *Given f, ϕ, and σ from the previous,*

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

Remark. The Brand classifying space $B(G,C)$ exists for any G (not necessarily finite) and any union of conjugacy classes (not just 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

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

We also define

$$M(G,C) \stackrel{\mathrm{def}}{=} H_2(B(G,C))_{\mathrm{tor}}.$$

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

$$f : \pi_1(S^2 \setminus [2k]) \to G$$

is an element of $\hat{R}_k$, with $k = a = b$. Then

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

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

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

By construction, the function

$$\mathrm{sch} : \hat{R}_k \to M(G,C)$$

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

Theorem 4.2 (Ellenberg-Venkatesh-Westerland [27, Thm. 1.1])**.** *Let G be a finite group generated by a conjugacy class C. Then, eventually, the Schur invariant yields a bijection*

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

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

Remark. Theorem 4.2 first appeared in version 1 of an arXiv e-print by Ellenberg, Venkatesh, and Westerland [12]. 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 [26]. The second author [27] then found a topological proof of the same result using the Brand classifying space. The new results of [27] also hold for surfaces with either genus or punctures or both, and thus subsume a result of Dunfield and Thurston [11, Thm. 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 [15].

We will use two basic properties of the Schur invariant, one of which requires a definition: Say that $f \in \hat{R}_k$ *bounds a plat* if there is an inclusion $S^2 \setminus [2k] \to B^3$ 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 $f \in \hat{R}_a$ and $g \in \hat{R}_b$, and $f \# g \in \hat{R}_{a+b}$ is their boundary sum, then* $\mathrm{sch}(f \# g) = \mathrm{sch}(f) + \mathrm{sch}(g)$.

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

Proof. Part 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:

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

Part 2: The map f is null-concordant by hypothesis, hence $\mathrm{sch}(f) = 0$ is null-homologous by Proposition 4.1. ☐

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

$$\hat{R}_k^0 \stackrel{\mathrm{def}}{=} \{f \in \hat{R}_k \mid \mathrm{sch}(f) = 0\}$$
$$R_k^0 \stackrel{\mathrm{def}}{=} \{f \in R_k \mid \mathrm{sch}(f) = 0\}.$$

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 [26, Sec. 4B]. If G is a perfect group, then it has a canonical central extension

$$M(G) \hookrightarrow \hat{G} \twoheadrightarrow G$$

called the *Schur cover* $\hat{G}$ of G. In general the conjugacy classes in $\hat{G}$ can be larger than their counterparts in G, in the sense that two preimages $g_1, g_2 \in \hat{G}$ 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

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

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

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

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*

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

We note Dunfield and Thurston proved a version of this lemma for maps from fundamental groups of closed surfaces, instead of punctured disks. Our limits are analogs of [11, Lems. 6.10 & 6.13].

Proof. For the first claim, we 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

$$g_k \stackrel{\text{def}}{=} c_{2k}^{-1} c_{2k-1} \ldots c_3 c_2^{-1} c_1$$

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 non-trivial equivalence relation $\sim$ such that M descends to a permutation on the quotient set $G/\sim$. Since the diagonal of M is entirely non-zero, 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, except with a numerator of $\#\hat{R}_k$ rather than $\#R_k$. For the rest of the limit, observe that in the given random process, the image of f_k is monotonic, more precisely that it almost surely increases to G and stays there. By contrast the condition $g_k = 1$ is recurrent. Therefore the limiting probability that $f_k \in \hat{R}_k$ is the same as the limiting probability that $f_k \in R_k$.

We reduce the second limit to the first one. Recall that $\tilde{G}$ is the central extension of G by $M(G,C)$, and that $\tilde{G}$ is a perfect group because the full Schur cover $\hat{G}$ of the perfect group G is perfect. Let $C' \subseteq \tilde{G}$ be a conjugacy class that lifts C. (Any such lift generates $\tilde{G}$.) Then $f \in R_k$ has a lift $f' \in R_k(\tilde{G}, C')$, and 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 $\tilde{G}$, 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

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

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

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

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

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

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*

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

Using part 1, this equivalence also commutes with the Schur invariant:

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

4. *The induced map*

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

is injective.

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

Part 2: Proposition 4.1, part 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.

Part 3: The main step is to review the construction of a map

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

representing $f \in \hat{R}_k(G)$, and to then relate ϕ to the map ψ, as promised in Proposition 4.1. Given

$$f: \pi_1(S^2 \setminus [2k]) \to G,$$

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

$$\phi: S^2 \setminus [2k] \to B(G)$$

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 ϕ because the same extension disk in S^2 is used ℓ times for each puncture. Moreover, if ψ_i is the ith component of the map ψ, then the composition $\phi_i = \psi_i \circ \phi$ matches the ith component f_i of f. Together with the proof of the Künneth formula and its use in part 1, this establishes that ith component of $\mathrm{sch}(f)$ is $\mathrm{sch}(f_i)$, as desired.

Part 4: Let $c \in C$ and consider the diagram

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

where each row is taken from part 4 of Proposition 4.1 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 [27]. Given any group homomorphism $f: A \times B \to G$, there is always a bi-additive 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 part 5 of Proposition 4.1. 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 ultra transitivity theorem

Theorem 4.2 says that the action of $B_{k,k}$ on R^0_k 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-Venkatesh [26, Thm. 5.1]. However, Theorem 4.7 refines this special case of Roberts-Venkatesh in the same way that our prior result [21, Thm. 5.1] refines a result of Dunfield-Thurston [11, Thm. 7.4].

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

$$U \stackrel{\text{def}}{=} \mathrm{Aut}(G,c) \subseteq \hat{U} \stackrel{\text{def}}{=} \mathrm{Aut}(G,C).$$

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

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

with an element $\alpha \in \mathrm{Aut}(G,C)$. It acts freely on the subset R^0_k because these homomorphisms are surjective. Moreover, the actions of $B_{k,k}$ and $\mathrm{Aut}(G,C)$ on R^0_k commute. In other words, the corresponding permutation representation is a map

$$\rho: B_{k,k} \to \mathrm{Sym}_{\hat{U}}(R^0_k).$$

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 $\hat{U}$-set ℓ-transitively on R_k^0.*

Proof. We choose k large enough so 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 $\hat{U}$-orbits and consider the product homomorphism

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

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

$$\sigma_{J,E} : B_{k,k} \to \mathrm{Sym}(T(J,E))$$

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

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

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 $\rho \times \sigma \times \phi$:*

$$\rho \times \phi \times \prod_{\substack{J \supseteq E \\ \#E < \#C}} \sigma_{J,E} : B_{k,k} \longrightarrow \mathrm{Sym}_{\hat{U}}(R_k^0) \times \mathrm{Sym}(k)^2 \times \prod_{\substack{J \supseteq E \\ \#E < \#C}} \mathrm{Sym}(T(J,E))$$

eventually contains $\mathrm{Rub}_{\hat{U}}(R_k^0)$, 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 $T(J,E)$ with $\#E < \#C$. It follows that such braids also act trivially on the set $\hat{R}_k \smallsetminus R_k$ of non-surjective maps in $\hat{R}_k$.

Proof. Following Dunfield-Thurston [11] and Roberts-Venkatesh [26], 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 $\mathrm{Rub}_{\hat{U}}(R_k^0)$.

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

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

By Lemmas 2.6 and 2.4, it suffices to show that $\mathrm{Alt}(R_k^0/\hat{U})$ is not a subquotient of $\mathrm{Sym}(T(J,E))$ 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 $\mathrm{Alt}(R_k^0/\hat{U})$ is eventually larger than any of these other groups. This follows from comparing $\#T(J,E) = \#E^{2k}$ to the bound in Lemma 4.4 that 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 [21] is a result of Dunfield and Thurston [11, Thm. 7.4] that they argue in the same way. However, they point out that they could use a non-classification result of Dixon and Mortimer [10, Thm. 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 [16], 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 non-abelian simple group, let $C \subseteq G$ be a non-trivial conjugacy class (which necessarily generates G), and let $c \in C$ be a distinguished element. Again, let

$$U \stackrel{\text{def}}{=} \mathrm{Aut}(G,c) \subseteq \mathrm{Aut}(G,C) \stackrel{\text{def}}{=} \hat{U}.$$

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 [21], we will face a new technical difficulty. Namely, even though the available braid actions are $\hat{U}$-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 $\hat{U}$-invariant choice. If we represented z using all of C or in any other $\hat{U}$-invariant manner, then the construction could only produce an intractable link invariant rather than specifically a knot invariant.

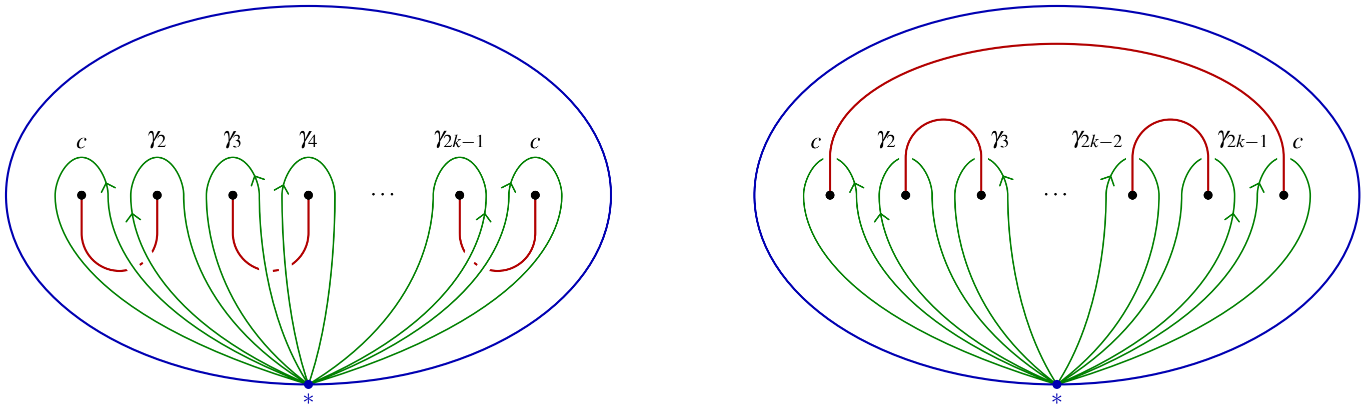

Figure 3. These two sets of red plats impose the initial and final constraints, respectively.

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 that we will later use to replace gates in a ZSAT circuit. Here a *gadget* is a semi-rigorous 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

$$z \stackrel{\text{def}}{=} (c,c,\dots,c) \in \hat{R}^0_k,$$

and let the alphabet be

$$A \stackrel{\text{def}}{=} \{z\} \cup \{(g_1,g_2,\dots,g_{2k}) \in R^0_k \mid g_1 = g_{2k} = c\}.$$

I.e., the non-zombie symbols in A are surjections with trivial Schur invariant such that the first and last punctures map to c specifically. The initialization and finalization 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

$$I \stackrel{\text{def}}{=} \{(g_1,g_2,\dots,g_{2k}) \in A \mid g_{2i} = g_{2i-1}\ \forall i \le k\}$$
$$F \stackrel{\text{def}}{=} \{(g_1,g_2,\dots,g_{2k}) \in A \mid g_{2i} = g_{2i+1}\ \forall i \le k-1\}$$

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 $D^2 \setminus [4k]$ be a pointed disk with $4k$ punctures, and divide it into two half-disks by a straight line, so that each half contains the base point 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

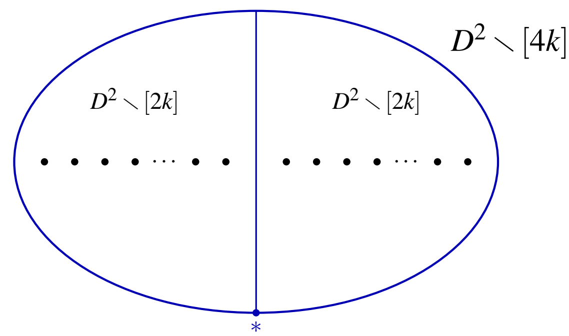

Figure 4. Splitting $D^2 \setminus [4k]$ as a boundary sum of two copies of $D^2 \setminus [2k]$.

takes $R^0_k \times R^0_k$ to a subset of R^0_{2k}. In particular, we identify A^2 with a subset of $R^0_{2k} \cup \{(z,z)\}$.

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 $R^0_{2k} \setminus \hat{U} \cdot (A^2)$.*
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,* i.e., *$b(\delta) \in PB_{4k} \le B_{2k,2k}$.*

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 $\hat{U}$-equivariant while A^2 is only closed under the action of U. This action has a unique $\hat{U}$-equivariant extension to $\hat{U} \cdot (A^2)$. Meanwhile property 3 implies that $b(\delta)$ acts trivially on $\hat{R}_{2k} \setminus R_{2k}$, so together the first three properties specify all of the action of $b(\delta)$ on $\hat{R}^0_{2k}$. However, $b(\delta)$ is not fully specified on all of $\hat{R}_{2k}$, because we place no restrictions on its effect on $f \in R^s_{2k}$ with non-vanishing 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 $\hat{R}_{2k}$ preserves all of the sets*

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

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

5.2. The reduction

Let Z be an instance of $\mathsf{ZSAT}_{U,A,I,F}$, with U,A,I,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.

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

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 $(D^2 \setminus [2kn])^{(m)}$ be a disk transverse to the braid, so that these disks and the braid gadgets alternate. Each disk $(D^2 \setminus [2kn])^{(m)}$ is also divided into subdisks $\{(D^2 \setminus [2kn])^{(m)}_i\}$ as before, with loops $\{\gamma^{(m)}_{j,i}\}$. Finally let $\gamma_0 \in \pi_1(S^3 \setminus K(Z))$ be the indicated meridian. In other words, $\gamma_0 = \gamma^{(0)}_{1,1}$.

We are interested in homomorphisms

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

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:

$$f^{(m)} : \pi_1((D^2 \setminus [2kn])^{(m)}) \to G$$
$$f^{(m)}_i : \pi_1((D^2 \setminus [2k])^{(m)}_i) \to G$$
$$f^{(m)}_{j,i} = f(\gamma^{(m)}_{j,i}) \in C.$$

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

$$p \stackrel{\text{def}}{=} f^{(0)} \qquad q \stackrel{\text{def}}{=} f^{(\ell)}.$$

The inclusion map

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

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, that there are no non-trivial 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 non-cyclic 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. Part 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.

Part 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 $(f^{(m-1)}_i, f^{(m-1)}_{i+1})$, 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.

Part 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,

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

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 $f \in H(K(Z),\gamma_0;G,c) \setminus X(Z)$ be a hypothetical spurious homomorphism. Then tautologically $p = f^{(0)} \in T^n_k \cong (C^{2k})^n$, but we quickly obtain an important restriction. Each p_i factors through the initial plat attached to

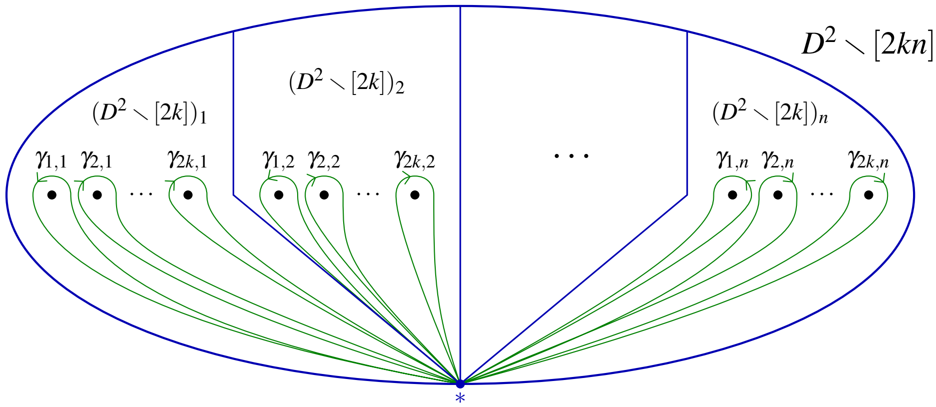

Figure 5. Punctured disks that encode n symbols of a ZSAT circuit.

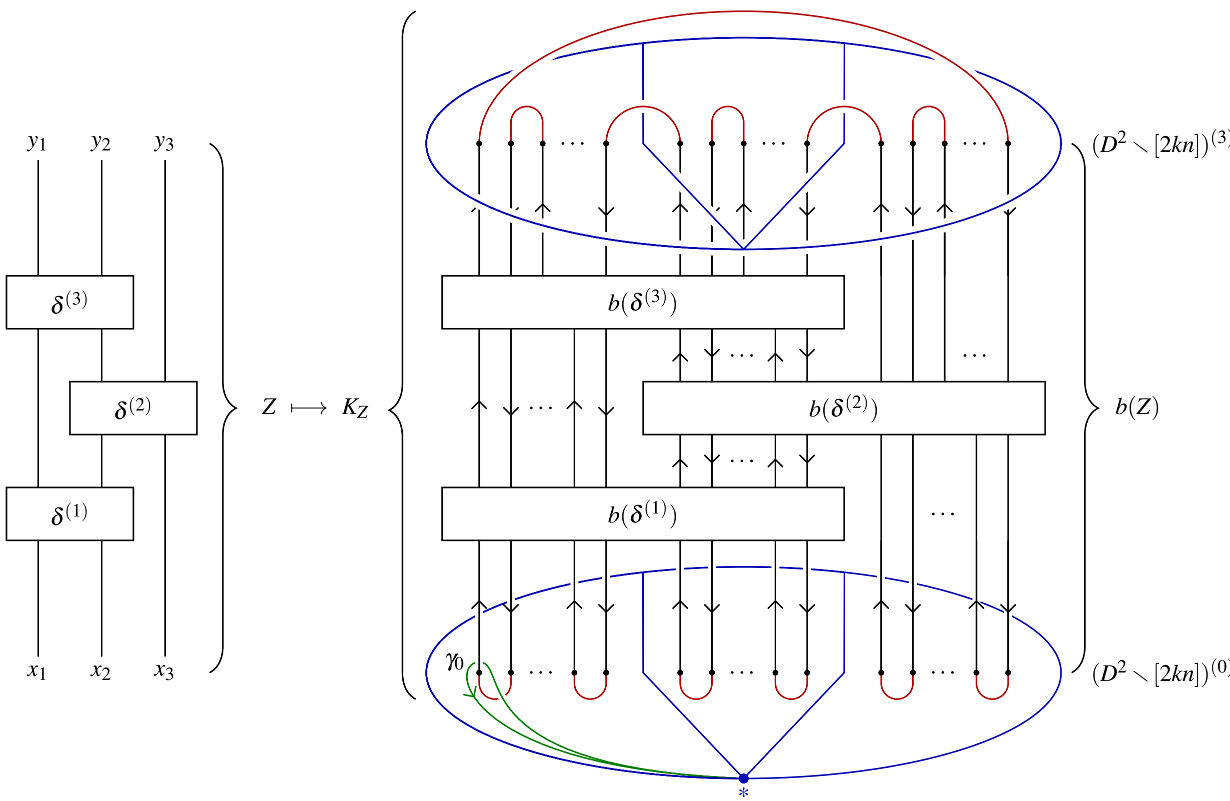

Figure 6. Reducing a circuit Z with $n = 3$ variables to the knot K_Z.

$(D^2 \setminus [2k])_i^{(0)}$, so Lemma 4.3 tells us that $\mathrm{sch}(p_i) = 0$ and thus that $p_i \in \hat{R}_k^0$. Moreover, Lemma 5.2 tells us that every braid gadget preserves this condition, so $f_i^{(m)} \in \hat{R}_k^0$ for every m and i. In other words, we can interpret $b(Z)$ as a circuit that uses the larger alphabet $\hat{R}_k^0 \supseteq A$.

We claim that we can further restrict the alphabet to $\hat{U} \cdot A$. If $p_i = f_i^{(0)} \in \mathbb{R}_k^0 \setminus \hat{U} \cdot A$ 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

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

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

We now show that $f^{(m)} \in \hat{U} \cdot A^n$ for every m. The condition that $f_i^{(m)} \in \hat{U} \cdot A$ 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 $p = f^{(0)} \in \hat{U} \cdot A^n$ 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 $b(\delta)$. Thus it is a parsimonious reduction from $\#\mathsf{ZSAT}_{U,A,I,F}$ to $\#H(-,G,c)$ that preserves the #P-completeness properties state in Theorem 3.1.

Remark. As in our previous work [21], the proof of Theorem 1.1 establishes an efficient bijection between $Q(K(Z),\gamma_0;G,c)$ and the orbits of non-trivial solutions to $\#\mathsf{ZSAT}_{U,A,I,F}$, and therefore the set of certificates in any problem in #P. This is an even stronger property than parsimonious reduction known as Levin reduction.

6. OPEN PROBLEMS

As with our previous theorem about homology 3-spheres [21], 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 [13], 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 [6].

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 non-abelian simple group.

---

[1] Ian Agol, Joel Hass, and William Thurston, *The computational complexity of knot genus and spanning area*, Trans. Amer. Math. Soc. **358** (2006), no. 9, 3821–3850, arXiv:math/0205057. 1.1

[2] Sanjeev Arora and Boaz Barak, *Computational complexity: a modern approach*, Cambridge University Press, Cambridge, 2009. 1.1

[3] Régis Barbanchon, *On unique graph 3-colorability and parsimonious reductions in the plane*, Theoret. Comput. Sci. **319** (2004), no. 1-3, 455–482. 1.1

[4] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani, *Strengths and weaknesses of quantum computing*, SIAM J. Comput. **26** (1997), no. 5, 1510–1523, arXiv:quant-ph/9701001. 1.1

[5] Neal Brand, *Classifying spaces for branched coverings*, Indiana Univ. Math. J. **29** (1980), no. 2, 229–248. 4.2

[6] Jin-Yi Cai, Aduri Pavan, and D. Sivakumar, *On the hardness of permanent*, STACS 99 (Trier), Lecture Notes in Comput. Sci., vol. 1563, Springer, Berlin, 1999, pp. 90–99. 6

[7] Richard H. Crowell and Ralph H. Fox, *Introduction to knot theory*, Graduate Texts in Mathematics, vol. 57, Springer-Verlag, 1977, Reprint of the 1963 original. 1

[8] Pierre de la Harpe and Vaughan F. R. Jones, *Graph invariants related to statistical mechanical models: examples and problems*, J. Combin. Theory Ser. B **57** (1993), no. 2, 207–227. 1

[9] Arnaud de Mesmay, Yo'av Rieck, Erick Sedgwick, and Martin Tancer, *The unbearable hardness of unknotting*, 2018, arXiv:1810.03502. 1.1

[10] John D. Dixon and Brian Mortimer, *Permutation groups*, Graduate Texts in Mathematics, vol. 163, Springer-Verlag, 1996. 4.4

[11] Nathan M. Dunfield and William P. Thurston, *Finite covers of random 3-manifolds*, Invent. Math. **166** (2006), no. 3, 457–521, arXiv:math/0502567. 2.1, 4.2, 4.3, 4.4, 4.4

[12] Jordan S. Ellenberg, Akshay Venkatesh, and Craig Westerland, *Homological stability for Hurwitz spaces and the Cohen-Lenstra conjecture over function fields, II*, 2012, arXiv:1212.0923. 1.2, 4.2, 4.2

[13] Chaim Even-Zohar, *Models of random knots*, J. Appl. Comput. Topol. **1** (2017), no. 2, 263–296, arXiv:1711.10470. 6

[14] Ralph H. Fox, *A quick trip through knot theory*, Topology of 3-manifolds and related topics (Proc. The Univ. of Georgia Institute, 1961), Prentice-Hall, 1962, pp. 120–167. 1

[15] Michael D. Fried and Helmut Völklein, *The inverse Galois problem and rational points on moduli spaces*, Math. Ann. **290** (1991), no. 4, 771–800. 1.2, 4.2

[16] Robert Gilman, *Finite quotients of the automorphism group of a free group*, Canad. J. Math **29** (1977), 541–551. 4.4

[17] Édouard Goursat, *Sur les substitutions orthogonales et les divisions régulières de l'espace*, Ann. Sci. École Norm. Sup. **6** (1889), 9–102. 2.1

[18] Philip Hall, *The Eulerian functions of a group*, Quart. J. Math. **7** (1936), 134–151. 2.1

[19] Mark Jerrum, Alistair Sinclair, and Eric Vigoda, *A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries*, J. ACM **51** (2004), no. 4, 671–697. 1.1

[20] Hari Krovi and Alexander Russell, *Quantum Fourier transforms and the complexity of link invariants for quantum doubles of finite groups*, Comm. Math. Phys. **334** (2015), no. 2, 743–777, arXiv:1210.1550. 1.1

[21] Greg Kuperberg and Eric Samperton, *Computational complexity and 3-manifolds and zombies*, arXiv:1707.03811, To appear in *Geom. Topol.* 1.1, 1.1, 1.2, 2.1, 2.2, 2.4, 2.2, 2.5, 2.6, 3, 3.1, 4.4, 4.4, 5, 5.2, 6

[22] James Munkres, *Algorithms for the assignment and transportation problems*, J. Soc. Indust. Appl. Math. **5** (1957), 32–38. 1.1

[23] Kurt Reidemeister, *Knoten und gruppen*, Abh. Math. Sem. Univ. Hamburg **5** (1927), no. 1, 7–23. 1

[24] Kenneth A. Ribet, *On l-adic representations attached to modular forms*, Invent. Math. **28** (1975), 245–275. 2.1

[25] ———, *Galois action on division points of Abelian varieties with real multiplications*, Amer. J. Math. **98** (1976), no. 3, 751–804. 2.1, 2.1, 2.3

[26] David P. Roberts and Akshay Venkatesh, *Hurwitz monodromy and full number fields*, Algebra Number Theory **9** (2015), no. 3, 511–545, arXiv:1401.7379. 4, 4.2, 4.2, 4.3, 4.4, 4.4

[27] Eric Samperton, *Schur-type invariants of branched G-covers of surfaces*, 2017, arXiv:1709.03182. 1.2, 4.2, 4.1, 4.2, 4.3

[28] L. G. Valiant and V. V. Vazirani, *NP is as easy as detecting unique solutions*, Theoret. Comput. Sci. **47** (1986), no. 1, 85–93. 1.1

[29] Leslie G. Valiant, *The complexity of computing the permanent*, Theoret. Comput. Sci. **8** (1979), no. 2, 189–201. 1.1

[30] *The Complexity Zoo*, `http://www.complexityzoo.com/`. 1.1