The Unbearable Hardness of Unknotting

Arnaud de Mesmay
Univ. Grenoble Alpes, CNRS, Grenoble INP[1], GIPSA-lab, 38000 Grenoble, France
arnaud.de-mesmay@gipsa-lab.fr

Yo'av Rieck
Department of Mathematical Sciences, University of Arkansas Fayetteville, AR 72701, USA
yoav@uark.edu

Eric Sedgwick
School of Computing, DePaul University, 243 S. Wabash Ave, Chicago, IL 60604, USA
ESedgwick@cdm.depaul.edu

Martin Tancer
Department of Applied Mathematics, Charles University,
Malostranské nám. 25, 118 00 Praha 1, Czech Republic
tancer@kam.mff.cuni.cz

Abstract

We prove that deciding if a diagram of the unknot can be untangled using at most k Reidemeister moves (where k is part of the input) is **NP**-hard. We also prove that several natural questions regarding links in the 3-sphere are **NP**-hard, including detecting whether a link contains a trivial sublink with n components, computing the unlinking number of a link, and computing a variety of link invariants related to four-dimensional topology (such as the 4-ball Euler characteristic, the slicing number, and the 4-dimensional clasp number).

2012 ACM Subject Classification Mathematics of computing → Geometric topology; Theory of computation → Problems, reductions and completeness

Keywords and phrases Knot, Link, NP-hard, Reidemeister move, Unknot recognition, Unlinking number, intermediate invariants

Digital Object Identifier 10.4230/LIPIcs.SoCG.2019.49

Related Version A full version of this paper is available at `https://arxiv.org/abs/1810.03502`.

Funding *Arnaud de Mesmay*: Partially supported by the French ANR projects ANR-16-CE40-0009-01 (GATO), the CNRS PEPS project COMP3D and the Czech-French collaboration project EMBEDS II (CZ: 7AMB17FR029, FR: 38087RM).
Yo'av Rieck: This work was partially supported by a grant from the Simons Foundation (grant number 283495 to Yo'av Rieck).
Martin Tancer: Partially supported by the Czech-French collaboration project EMBEDS II (CZ: 7AMB17FR029, FR: 38087RM), by the project CE-ITI (GAČR P202/12/G061) and by the Charles University projects PRIMUS/17/SCI/3 and UNCE/SCI/004.

Acknowledgements We thank Amey Kaloti and Jeremy Van Horn Morris for many helpful conversations. We are grateful to the anonymous referees for a careful reading of the paper and many helpful suggestions.

[1] Institute of Engineering Univ. Grenoble Alpes

35th International Symposium on Computational Geometry (SoCG 2019).
Editors: Gill Barequet and Yusu Wang; Article No. 49; pp. 49:1–49:19
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Unknot recognition via Reidemeister moves. The *unknot recognition problem* asks whether a given knot is the unknot. Decidability of this problem was established by Haken [10], and since then several other algorithms were constructed (see, e.g., the survey of Lackenby [20]).

One can ask, naively, if one can decide whether a given knot diagram represents the unknot simply by untangling the diagram: trying various Reidemeister moves until there are no more crossings. The problem is knowing when to stop: if we have not been able to untangle the diagram using so many moves, is the knot in question necessarily knotted or should we keep on trying? "Hard unknots" can be found, for example, in [14].

In [11], Hass and Lagarias gave an explicit (albeit rather large) bound on the number of Reidemeister moves needed to untangle a diagram of the unknot. Lackenby [18] improved the bound to polynomial thus showing that the unknot recognition problem is in **NP** (this was previously proved in [12]). The unknot recognition problem is also in co-**NP** [19] (assuming the Generalized Riemann Hypothesis, this was previously shown by Kuperberg [16]). Thus if the unknot recognition problem were **NP**-complete (or co-**NP**-complete) we would have that **NP** and co-**NP** coincide which is commonly believed not to be the case; see, for example, [9, Section 7.1]. This suggests that the unknot recognition problem is *not* **NP**-hard.

It is therefore natural to ask if there is a way to use Reidemeister moves leading to a better solution than a generic brute-force search. Our main result suggests that there may be serious difficulties in such an approach: given a 3-SAT instance Φ we construct an unknot diagram and a number k, so that the diagram can be untangled using at most k Reidemeister moves if and only if Φ is satisfiable. Hence any algorithm that can calculate the minimal number of Reidemeister moves needed to untangle unknot diagrams will be robust enough to tackle any problem in **NP**:

▶ **Theorem 1.** *Given an unknot diagram D and an integer k, deciding if D can be untangled using at most k Reidemeister moves is* **NP***-complete.*

NP-membership follows from the result of Lackenby [18], so we only show **NP**-hardness. For the reduction in the proof of Theorem 1 we have to construct arbitrarily large diagrams of the unknot. The difficulty in the proof is to establish tools powerful enough to provide useful lower bounds on the minimal number of Reidemeister moves needed to untangle these diagrams. For instance, the algebraic methods of Hass and Nowik [13] do not appear strong enough for our reduction. It is also quite easy to modify the construction and give more easily lower bounds on the number of Reidemeister moves needed to untangle unlinks if one allows the use of arbitrarily many components of diagrams with constant size, but those techniques too cannot be used for Theorem 1. We develop the necessary tools in Section 4.

Computational problems for links. Our approach for proving Theorem 1 partially builds on techniques to encode satisfiability instances using Hopf links and Borromean rings, that we previously used in [7] (though the technical details are very different). With these techniques, we also show that a variety of link invariants are **NP**-hard to compute. Precisely, we prove:

▶ **Theorem 2.** *Given a link diagram L and an integer k, the following problems are* **NP***-hard:*
(a) *deciding whether L admits a trivial unlink with k components as a sublink.*
(b) *deciding whether an intermediate invariant has value k on L,*
(c) *deciding whether $\chi_4(L) = 0$,*
(d) *deciding whether L admits a smoothly slice sublink with k components.*

We refer to Definition 13 for the definitions of $\chi_4(L)$, the 4-ball Euler characteristic, and intermediate invariants. These are broadly related to the topology of the 4-ball, and include the unlinking number, the ribbon number, the slicing number, the concordance unlinking number, the concordance ribbon number, the concordance slicing number, and the 4-dimensional clasp number. See, e.g., [25] for a discussion of many intermediate invariants.

Related results. The complexity of problems with knots and links is poorly understood. In particular, only very few computational lower bounds are known, and, as far as we know, almost none concern classical knots (i.e., knots embedded in S^3): apart from our Theorem 1, the only other such hardness proof we know of [17, 24] concerns counting coloring invariants of knots. More lower bounds are known for classical links. Lackenby [21] showed that determining if a link is a sublink of another one is **NP**-hard. Our results strengthen this by showing that even finding a trivial sublink is already **NP**-hard. Agol, Hass and Thurston [1] showed that computing the genus of a knot in a 3-manifold is **NP**-hard, and Lackenby [21] showed that computing the Thurston complexity of a link in S^3 is also **NP**-hard. Our results complement this by showing that the 4-dimensional version of this problem is also **NP**-hard.

Regarding upper bounds, the current state of knowledge is only slightly better. While, as we mentioned before, it is now known that the unknot recognition problem is in **NP**$\cap$ co-**NP**, many natural link invariants are not even known to be decidable. In particular, this is the case for all the invariants for which we prove **NP**-hardness, except for the problem of finding the maximal number of components of a trivial sublink, which is in **NP** (see Theorem 5).

Shortly before we finished our manuscript, Koenig and Tsvietkova posted a preprint [15] that also shows that certain computational problems on links are **NP**-hard, with some overlap with the results obtained in this paper (the trivial sublink problem and the unlinking number). They also show **NP**-hardness of computing the number of Reidemeister moves for turning one unlink diagram into another, but their construction does not untangle the diagram and requires arbitrarily many components. Theorem 1 of the current paper is stronger and answers Question 17 of [15].

Organization. After some preliminaries in Section 2, we start by sketching the hardness of the trivial sublink problem in Section 3 because it is very simple and provides a good introduction for our other reductions. We then proceed to prove Theorem 1 in Sections 4 and 5 and the hardness of the unlinking number and the other invariants in Sections 6 and 7. These three proofs are mostly independent and the reader can read any of them separately.

2 Preliminaries

Notation. Most of the notation we use is standard. By *knot* we mean a tame piecewise linear embedding of the circle S^1 into the 3-sphere S^3. By *link* we mean a tame, piecewise linear embedding of the disjoint union of any finite number of copies of S^1. We assume basic familiarity with computational complexity and knot theory, and refer to basic textbooks such as Arora and Barak [5] for the former and Rolfsen [23] for the latter.

Diagram of a knot or a link. All the computational problems that we study in this paper take as input the *diagram* of a knot or a link, which we define here. A diagram of a knot is a piecewise linear map $D\colon S^1 \to \mathbb{R}^2$ in general position, obtained by composing the embedding $S^1 \to \mathbb{R}^3$ with a projection $\mathbb{R}^3 \to \mathbb{R}^2$. For such a map, every point in $\mathbb{R}^2$ has at most two preimages, and there are finitely many points in $\mathbb{R}^2$ with exactly two preimages (called *crossings*); at each crossing we indicate which arc crosses over and which crosses under. We similarly define a link diagram; for details see, for example [23].

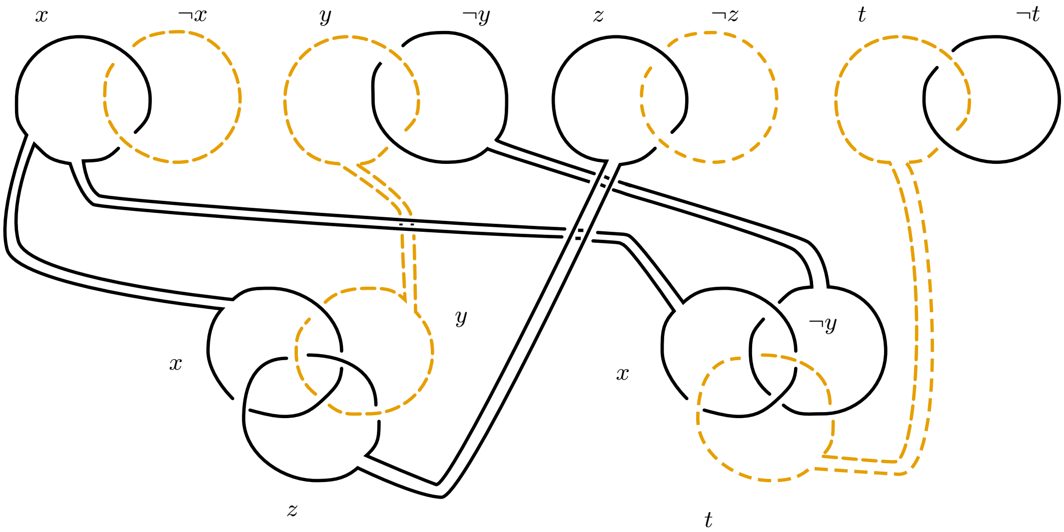

Figure 1 L_Φ for $\Phi = (x \vee y \vee z) \wedge (x \vee \neg y \vee t)$, a satisfying assignment is dashed.

By an *arc* in the diagram D we mean a set $D(\alpha)$ where α is an arc in S^1, i.e., a subset of S^1 homeomorphic to the closed interval (note that this definition is slightly non-standard).

The size of a knot or a link diagram is its number of crossings plus number of components of the link. Up to a constant factor, this complexity exactly describes the complexity of encoding the combinatorial information contained in a knot or link diagram.

3-satisfiability. In our reductions, we use the NP-hardness of the 3-SAT problem. An input is a formula in conjunctive normal form and we assume that each clause contains exactly three literals. For the proof of Theorem 1, we need a slightly restricted form of the 3-SAT problem given in the lemma below. The proof is simple (see Lemma 5 of the full version [8]).

▶ **Lemma 3** (Probably folklore). *Deciding whether a formula Φ in conjunctive normal form is satisfiable is* **NP***-hard even if we assume the following conditions on Φ.*
- *Each clause contains exactly three literals.*
- *No clause contains both x and $\neg x$ for some variable x.*
- *Each pair of literals $\{\ell_1, \ell_2\}$ occurs in at most one clause.*

3 Trivial sublink

Informally, the trivial sublink problem asks, given a link L and a positive integer n, whether L admits the n-component unlink as a sublink. We define:

▶ **Definition 4** (The Trivial Sublink Problem). *An* unlink, *or a* trivial link, *is a link in S^3 whose components bound disjointly embedded disks. A* trivial sublink *of a link L is an unlink formed by a subset of the components of L. The* trivial sublink problem *asks, given a link L and a positive integer n, whether L admits an n component trivial sublink.*

▶ **Theorem 5.** *The trivial sublink problem is* **NP***-complete.*

For a complete proof of this theorem see Theorem 4 of the full version [8], yet it is simple and intuitive enough to be explained in a few words and a picture here.

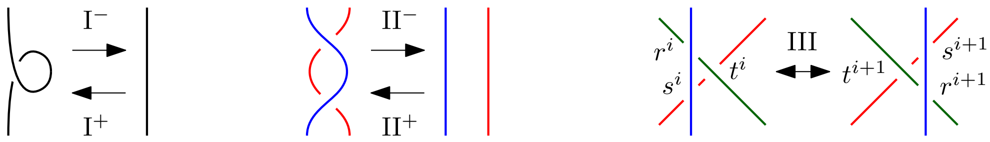

Figure 2 Reidemeister moves.

Sketch. **NP** membership follows from Hass, Lagarias and Pippenger [12] (also Lackenby [18]). For **NP**-hardness, starting with a 3-SAT instance Φ with n variables, we construct a $2n$-component link L_Φ (see Figure 1) consisting of a Hopf link for each variable and Borromean rings for each clause (each component corresponds to a literal as labeled). Each Borromean rings component is banded to the Hopf link component with the same label.

Given a satisfying assignment ($x = z =$ FALSE, $y = t =$ TRUE in Figure 1), remove the components of satisfied literals (dashed yellow). Then from each set of Borromean rings at least one ring was removed; thus the sublink retracts into n separated unknots.

Conversely, let U_n be an n-component trivial sublink of L_Φ (black in Figure 1). Since the Hopf link cannot be found in U_n, for each variable x, one of the components corresponding to x and $\neg x$ is in U_n, and the other is not. Since the Borromean rings cannot be found in U_n from each set Borromean rings at least one ring is *not* in U_n. We conclude that the components *not* in U_n define a satisfying assignment for Φ. ◀

4 The defect

Reidemeister moves. Reidemeister moves are local modifications of a diagram depicted in Figure 2 (the labels at the crossings in a III move will be used only later on). We distinguish the I move (left), the II move (middle) and the III move (right). The first two moves affect the number of crossings, thus we further distinguish the I^- and the II^- moves which reduce the number of crossings from the I^+ and the II^+ moves which increase the number of crossings.

The number of Reidemeister moves for untangling a knot. A diagram of the unknot is *untangled* if it does not contain any crossings. The untangled diagram is denoted by U. Given a diagram D of an unknot, an *untangling* of D is a sequence $\mathbf{D} = (D^0, \dots, D^k)$ where $D^0 = D$, $D^k = U$ (recall that diagrams are only considered up to isotopy) and D^i is obtained from D^{i-1} by a single Reidemeister move. The number of Reidemeister moves in $\mathbf{D}$ is denoted by $\mathrm{rm}(\mathbf{D})$, that is, $\mathrm{rm}(\mathbf{D}) = k$. We also define $\mathrm{rm}(D) := \min \mathrm{rm}(\mathbf{D})$ where the minimum is taken over all untanglings $\mathbf{D}$ of D.

The defect. Let us denote by $\mathrm{cross}(D)$ the number of crossings in D. We define the *defect* of an untangling $\mathbf{D}$ by the formula

$$\mathrm{def}(\mathbf{D}) := 2\,\mathrm{rm}(\mathbf{D}) - \mathrm{cross}(D).$$

The *defect* of a diagram D is defined as $\mathrm{def}(D) := 2\,\mathrm{rm}(D) - \mathrm{cross}(D)$. Equivalently, $\mathrm{def}(D) = \min \mathrm{def}(\mathbf{D})$ where the minimum is taken over all untanglings $\mathbf{D}$ of D. The defect is a convenient way to reparametrize $2\,\mathrm{rm}(\mathbf{D})$ due to the following observation.

▶ **Observation 6.** *For any diagram D of the unknot and any untangling $\mathbf{D}$ of D we have $\mathrm{def}(\mathbf{D}) \geq 0$. Equality holds if and only if $\mathbf{D}$ uses only II^- moves.*

This notion of defect is different from the one that was introduced by Chang and Erickson [6] (following Arnold [3, 4] and Aicardi [2]) to study homotopy moves.

Proof. A Reidemeister move in $\mathbf{D} = (D^0, \dots, D^k)$ removes at most two crossings and the II^- move is the only move that removes exactly two crossings. Thus, the number of crossings in $D = D^0$ is at most $2k$ and equality holds if and only if every move is a II^- move. ◀

Crossings contributing to the defect. Let $\mathbf{D} = (D^0, \dots, D^k)$ be an untangling of a diagram $D = D^0$ of an unknot. Given a crossing r^i in D^i, for $0 \le i \le k-1$, it may vanish by the move transforming D^i into D^{i+1} if this is a I^- or a II^- move affecting the crossing. In all other cases it *survives* and we denote by r^{i+1} the corresponding crossing in D_{i+1}. Note that in the case of a III move there are three crossings affected by the move and three crossings afterwards. Both before and after, each crossing is the unique intersection between a pair of the three arcs of the knot that appear in this portion of the diagram. So we may say that these three crossings survive the move though they change their actual geometric positions (they swap the order in which they occur along each of the three arcs); see Figure 2.

With a slight abuse of terminology, by a *crossing in* $\mathbf{D}$ we mean a maximal sequence $\mathbf{r} = (r^a, r^{a+1}, \dots, r^b)$ such that r^{i+1} is the crossing in D^{i+1} corresponding to r^i in D^i for any $a \le i \le b-1$. By maximality we mean that r^b vanishes after the $(b+1)$st move and either $a = 0$ or r^a is introduced by the ath Reidemeister move (which must be a I^+ or II^+ move).

An *initial crossing* is a crossing $\mathbf{r} = (r^0, r^1, \dots, r^b)$ in $\mathbf{D}$. Initial crossings in $\mathbf{D}$ are in one-to-one correspondence with crossings in $D = D^0$. For simplicity of notation, r^0 is also denoted r (as a crossing in D).

A Reidemeister II^- move in $\mathbf{D}$ is *economical* if both crossings removed by this move are initial crossings; otherwise, it is *wasteful*. Let $m_3(\mathbf{r})$ be the number of III moves affecting a crossing $\mathbf{r}$. The *weight* of an initial crossing $\mathbf{r}$ is defined in the following way.

$$w(\mathbf{r}) = \frac{2}{3} m_3(\mathbf{r}) + \begin{cases} 0 & \text{if } \mathbf{r} \text{ vanishes by an economical } \mathrm{II}^- \text{ move;} \\ 1 & \text{if } \mathbf{r} \text{ vanishes by a } \mathrm{I}^- \text{ move;} \\ 2 & \text{if } \mathbf{r} \text{ vanishes by a wasteful } \mathrm{II}^- \text{ move.} \end{cases}$$

For later purposes, we also define $w(r) := w(\mathbf{r})$ and $w(R) := \sum_{r \in R} w(r)$ for a subset R of the set of all crossings in D.

▶ **Lemma 7.** *Let* $\mathbf{D}$ *be an untangling of a diagram* D*. Then*

$$\operatorname{def}(\mathbf{D}) \ge \sum_{\mathbf{r}} w(\mathbf{r}),$$

where the sum is over all initial crossings $\mathbf{r}$ *of* $\mathbf{D}$*. If, in addition,* $\mathbf{D}$ *uses only the the* I^- *and* II^- *moves only, then (for the same sum) we get*

$$\operatorname{def}(\mathbf{D}) = \sum_{\mathbf{r}} w(\mathbf{r}) = \textit{ number of } \mathrm{I}^- \textit{ moves.}$$

The proof uses a discharging technique; see Lemmas 7 and 8 of the full version [8].

Twins and the preimage of a bigon. Let $\mathbf{r}$ be an initial crossing in an untangling $\mathbf{D} = (D^0, \dots, D^k)$ removed by an economical II^- move. The *twin* of $\mathbf{r}$, denoted by $t(\mathbf{r})$ is the other crossing in $\mathbf{D}$ removed by the same II^- move. Note that $t(\mathbf{r})$ is also an initial crossing (because the move is economical). We also get $t(t(\mathbf{r})) = \mathbf{r}$. If $\mathbf{r} = (r^0, \dots, r^b)$, then we extend the definition of a twin to D^i in such a way that $t(r^i)$ is uniquely defined by $t(\mathbf{r}) = (t(r^0), \dots, t(r^b))$. In particular, $t(r)$ is a twin of a crossing $r = r^0$ in D (if it exists).

Furthermore, the crossings r^b and $t(r^b)$ in D^b form a bigon that is removed by the forthcoming II^- move. Let $\alpha^b(r)$ and $\beta^b(r)$ be the two arcs of the bigon (with endpoints r^b and $t(r^b)$) so that $\alpha^b(r)$ is the arc that, after extending slightly, overpasses the crossings r^b and $t(r^b)$ whereas a slight extension of $\beta^b(r)$ underpasses these crossings. (The reader may remember this as α is "above" and β is "below".) Now we can inductively define arcs $\alpha^i(r)$ and $\beta^i(r)$ for $0 \leq i \leq b-1$ so that $\alpha^i(r)$ and $\beta^i(r)$ are the unique arcs between r^i and $t(r^i)$ which are transformed to (already defined) $\alpha^{i+1}(r)$ and $\beta^{i+1}(r)$ by the $(i+1)$st Reidemeister move. We also set $\alpha(r) = \alpha^0(r)$ and $\beta(r) = \beta^0(r)$. Intuitively, $\alpha(r)$ and $\beta(r)$ form a preimage of the bigon removed by the bth move and we call them the *preimage arcs* between r and $t(r)$.

Close neighbors. Let R be a subset of the set of crossings in D. Let r and s be any two crossings in D (not necessarily in R) and let $c \geq 0$ be an integer. We say that r and s are *c-close neighbors with respect to R* if r and s can be connected by two arcs α and β such that

- α enters r and s as an overpass and β enters r and s as an underpass;
- α and β may have self-crossings; however, neither r nor s is in the interior of α or β; and
- α and β together contain at most c crossings from R in their interiors. (If there is a crossing in the interior of both α and β, this crossing is counted only once.)

▶ **Lemma 8.** *Let R be a subset of the set of crossings in D, let $c \in \{0,1,2,3\}$. Let r be the crossing in R which is the first of the crossings in R removed by an economical II^- move (we allow a draw). If $w(R) \leq c$, then r and its twin $t(r)$ are c-close neighbors with respect to R.*

Sketch. See Lemma 9 of the full version [8]. Let $\alpha(r)$ and $\beta(r)$ be the preimage arcs between r and $t(r)$. We want to verify that they satisfy the properties of the arcs from the definition of the close neighbors. The first item follows immediately from the definition of preimage arcs. For the second item, if we had r or $t(r)$ in the interior of α or β then we cannot get a bigon between r^b and $t(r^b)$ by subsequent moves. The third item is verified by an analysis of Reidemeister moves removing the crossings of α and β before we get a bigon between r^b and $t(r^b)$, using $c \leq 3$: if r and $t(r)$ are not c-close with respect to R, then the Reidemeister moves needed to "clean" the bigon for the II^- move contribute more than c to $w(R)$. ◀

5 The reduction

Let Φ be a formula in conjunctive normal form satisfying the conditions stated in Lemma 3 and let n be the number of variables. Our aim is to build a diagram $D(\Phi)$ by a polynomial-time algorithm such that $\mathrm{def}(D(\Phi)) \leq n$ if and only if Φ is satisfiable.

The variable gadget. First we describe the variable gadget. For every variable x we consider the diagram depicted at Figure 3 and we denote it $V(x)$.

The gadget contains 17 crossings $p[x], p[\neg x], q[x], q[\neg x]$, $r(x)$ and $s_i(x)$ for $1 \leq i \leq 12$.

The variable gadget also contains six distinguished arcs $\gamma_i[x]$ and $\gamma_i[\neg x]$ (for $i = 1, 2$), $\delta(x)$ and $\varepsilon(x)$ and six distinguished auxiliary points $a_1(x), \dots, a_6(x)$ which will be useful later on in order to describe how the variable gadget is used in the diagram $D(\Phi)$.

We also call the arc between $a_1(x)$ and $a_2(x)$ which contains $\gamma_1[\neg x]$ and $\gamma_2[\neg x]$ the $\neg x$ *tentacle*, and similarly, we have removed $\neg x$ tentacle between $a_2(x)$ and $a_3(x)$. Informally, a satisfying assignment to Φ will correspond to the choice whether we will decide to remove first the loop at $p[x]$ by a I^- move and simplify the x tentacle or whether we remove first the loop at $p[\neg x]$ and remove the $\neg x$ tentacle in the final construction of $D(\Phi)$.

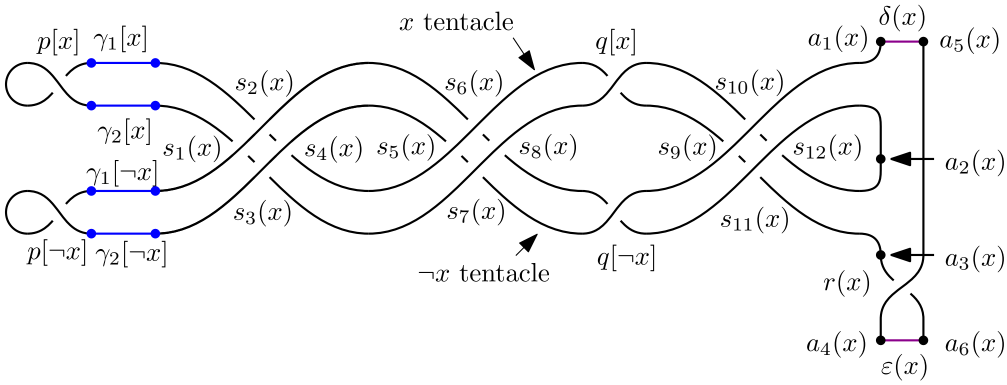

Figure 3 The variable gadget $V(x)$.

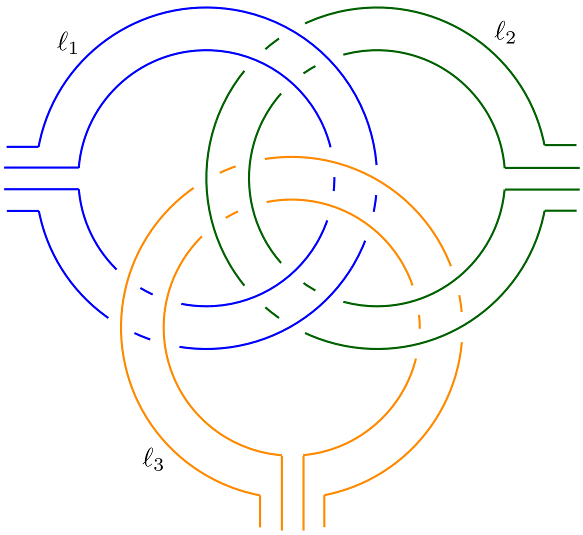

Figure 4 The clause gadget for clause $c = (\ell_1 \vee \ell_2 \vee \ell_3)$.

We also remark that in the notation, we use square brackets for objects that come in pairs and will correspond to a choice of literal $\ell \in \{x, \neg x\}$. This regards $p[\ell]$, $q[\ell]$, $\gamma_1[\ell]$ and $\gamma_2[\ell]$ whereas we use parentheses for the remaining objects.

The clause gadget. Given a clause $c = (\ell_1 \vee \ell_2 \vee \ell_3)$ in Φ, the *clause gadget* is depicted at Figure 4.The construction is based on the Borromean rings. It contains three pairs of arcs (distinguished by color) and with a slight abuse of notation, we refer to each of the three pairs of arcs as a "ring". Note that each ring has four pendent endpoints (or leaves) as in the picture. Each ring corresponds to one of the literals ℓ_1, ℓ_2, and ℓ_3.

A blueprint for the construction. Now we build a blueprint for the construction of $D(\Phi)$. Let $x_1, \dots, x_n$ be the variables of Φ and let $c_1, \dots, c_m$ be the clauses of Φ. For each clause $c_j = (\ell_1 \vee \ell_2 \vee \ell_3)$ we take a copy of the graph $K_{1,3}$ (also known as the star with three leaves). We label the vertices of degree 1 of such a $K_{1,3}$ by the literals ℓ_1, ℓ_2, and ℓ_3. Now we draw these stars into the plane sorted along a horizontal line; see Figure 5. Next for each literal $\ell \in \{x_1, \dots, x_n, \neg x_1, \dots, \neg x_n\}$ we draw a piecewise linear segment containing all vertices labelled with that literal according to the following rules (follow Figure 5).

$$\Phi = (x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee x_2 \vee \neg x_3) \wedge (\neg x_1 \vee \neg x_2 \vee x_4) \wedge (x_1 \vee \neg x_3 \vee \neg x_4)$$

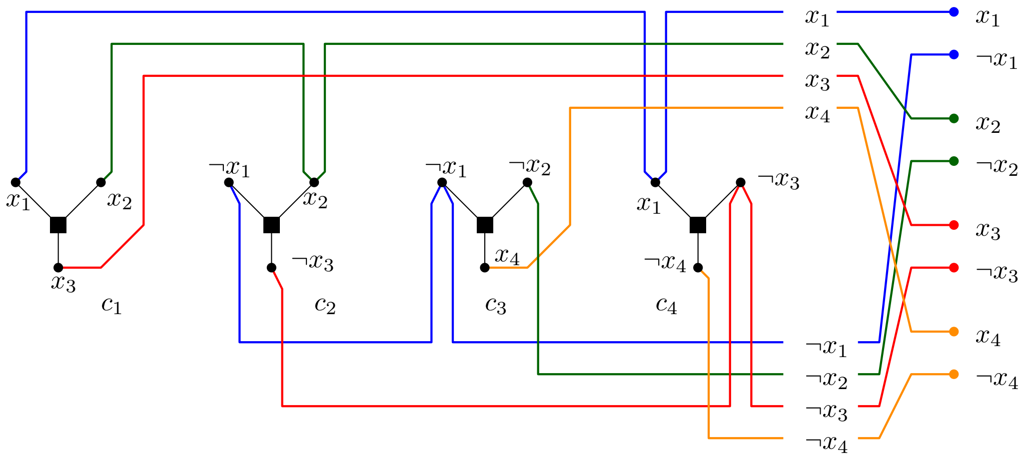

Figure 5 A blueprint for the construction of $D(\Phi)$.

- The segments start on the right of $K_{1,3}$'s in the top down order $x_1, \neg x_1, x_2, \ldots, x_n, \neg x_n$.
- They continue to the left while we permute them to the order $x_1, \ldots, x_n, \neg x_1, \ldots, \neg x_n$. We also require that $x_1, \ldots, x_n$ occur above the graphs $K_{1,3}$ and $\neg x_1, \ldots, \neg x_n$ occur below these graphs (everything is still on the right of the graphs).
- For each literal ℓ the segment for ℓ continues to the left while it makes a "detour" to each vertex v labelled ℓ. If v is not the leftmost vertex labelled ℓ, then the detour is done by a "finger" of two parallel lines. Each finger avoids $K_{1,3}$'s except of v. If v is the leftmost, then we perform only a half of the finger so that v is the endpoint of the segment.

Note that the segments often intersect each other; however, for any i the segments for x_i and $\neg x_i$ do not intersect (as we assume that no clause contains both x_i and $\neg x_i$).

The final diagram. Finally, we explain how to build the diagram $D(\Phi)$ from the blueprint.

Step I (four parallel segments): We replace each segment for a literal ℓ with four parallel segments; see Figure 6. The outer two will correspond to the arc $\gamma_1[\ell]$ from the variable gadget and the inner two will correspond to $\gamma_2[\ell]$; compare with Figure 3.

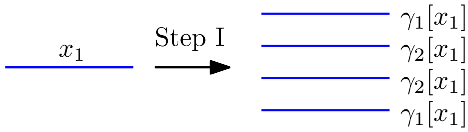

Figure 6 Step I: Replacing segments.

Step II (clause gadgets): We replace each copy of $K_{1,3}$ by a clause gadget for the corresponding clause c; see Figure 7. Now we aim to describe how is the clause gadget connected to the quadruples of parallel segments obtained in Step I. Let v be a degree 1 vertex of the $K_{1,3}$ we are just replacing. Let ℓ be the literal which is the label of this vertex. Then c may or may not be the leftmost clause containing a vertex labelled ℓ.
If c is the leftmost clause containing a vertex labelled ℓ, then there are four parallel segments for ℓ with pendent endpoints (close to the original position of v) obtained in

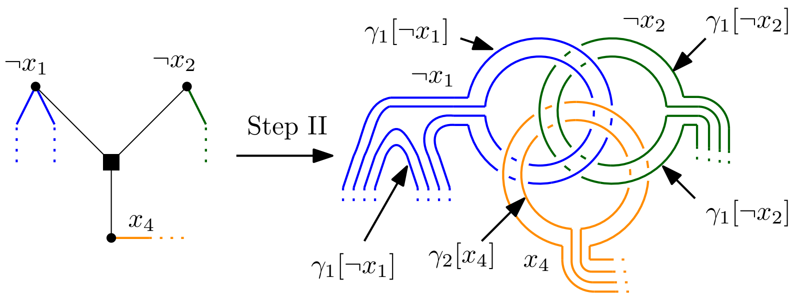

Figure 7 Step II: Replacing the $K_{1,3}$.

Step I. We connect them to the pendent endpoints of the clause gadget (on the ring for ℓ); see $\neg x_2$ and x_4 in Figure 7. Note also that at this moment the two $\gamma_1[\ell]$ arcs introduced in Step I merge as well as the two $\gamma_2[\ell]$ arcs merge.

If c is not the leftmost clause labelled ℓ then there are four parallel segments passing close to v (forming a tip of a finger from the blueprint). We disconnect the two segments closest to the tip of the finger and connect them to the pendent endpoints of the clause gadget (on the ring for ℓ); see $\neg x_1$ in Figure 7.

Step III (resolving crossings): If two segments in the blueprint, corresponding to literals ℓ and ℓ' cross, Step I blows up each such crossing into 16 crossing of corresponding quadruples. We resolve overpasses/underpasses at these crossings in the same way; see Figure 8.

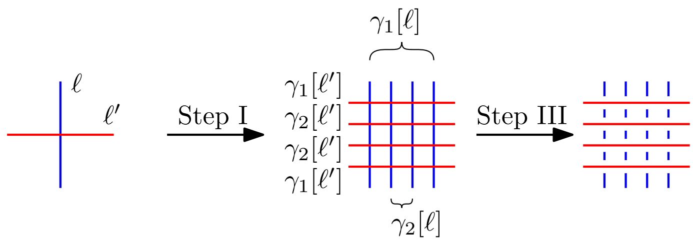

Figure 8 Step III: Resolving crossings.

However, we require one additional condition on the choice of overpasses/underpasses. If ℓ and ℓ' appear simultaneously in some clause c we have 8 crossings on the rings for ℓ and ℓ' in the clause gadget for c. We can assume that the ring of ℓ passes over the ring of ℓ' at all these crossings (otherwise we swap ℓ and ℓ'). Then for the 16 crossings on segments for ℓ and ℓ' we pick the other option, that is we want that the $\gamma_1[\ell]$ and $\gamma_2[\ell]$ arcs underpass the $\gamma_1[\ell']$ and $\gamma_2[\ell']$ arcs at these crossings. This is a globally consistent choice because we assume that there is at most one clause containing both ℓ and ℓ', this is the third condition in the statement of Lemma 3.

Step IV (the variable gadgets): For each variable x_i, the segments $\gamma_1[x_i], \gamma_2[x_i], \gamma_1[\neg x_i]$ and $\gamma_2[\neg x_i]$ do not intersect each other. We extend them to a variable gadget as in Figure 9. Namely, to the bottom right endpoints of $\gamma_1[x_i], \gamma_2[x_i], \gamma_1[\neg x_i]$ and $\gamma_2[\neg x_i]$ we glue the parts of the variable gadget containing the crossings $p[x_i]$ and $p[\neg x_i]$ and to the top right endpoints of $\gamma_1[x_i], \gamma_2[x_i], \gamma_1[\neg x_i]$ and $\gamma_2[\neg x_i]$ we glue the remainder of the variable gadget. At this moment, we obtain a diagram of a link, where each link component has a diagram isotopic to the diagram on Figure 3.

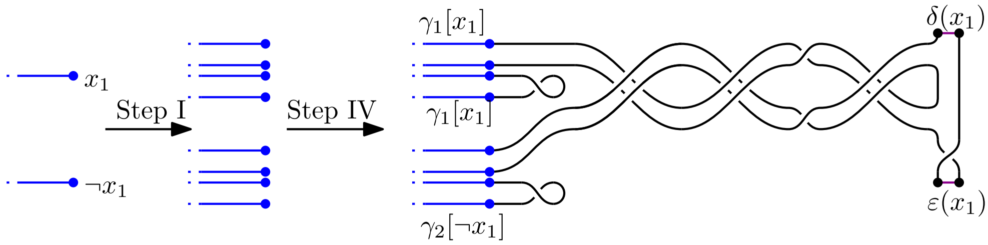

Figure 9 Step IV: Adding the variable gadgets.

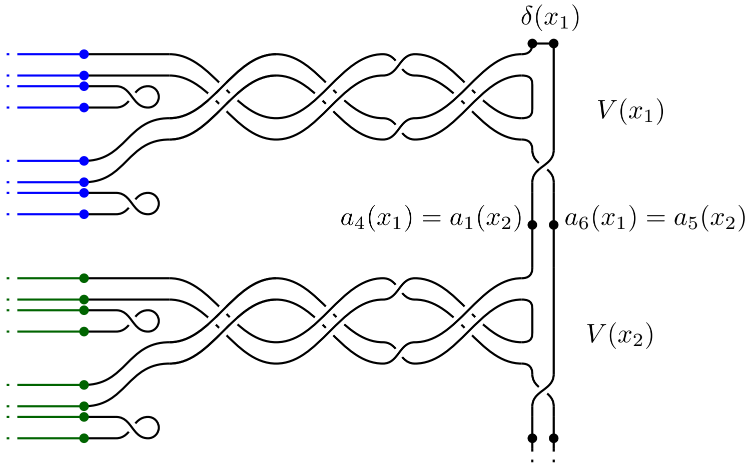

Figure 10 Step V: Interconnecting the variable gadgets.

Step V (interconnecting the variable gadgets): Finally, we form a connected sum of individual components. Namely, for every $1 \leq i \leq n-1$ we perform the knot sum along the arcs $\delta(x_i)$ and $\varepsilon(x_{i+1})$ by removing them and identifying $a_4(x_i)$ with $a_1(x_{i+1})$ and $a_6(x_i)$ with $a_5(x_{i+1})$ as on Figure 10. The arcs $\delta(x_1)$ and $\varepsilon(x_n)$ remain untouched. This way we obtain the desired unknot diagram $D(\Phi)$; see Figure 11.

The core of the **NP**-hardness reduction is the following theorem.

▶ **Theorem 9.** *Let Φ be a formula with n variables in conjunctive normal form satisfying the conditions in the statement of Lemma 3. Then* $\operatorname{def}(D(\Phi)) \leq n$ *if and only if Φ is satisfiable.*

Theorem 1 immediately follows from Theorem 9 and Lemma 3:

Proof of Theorem 1 modulo Theorem 9. Due to the definition of the defect, the minimum number of Reidemeister moves required to untangle D equals $\frac{1}{2}(\operatorname{def}(D(\Phi)) + \operatorname{cross}(D(\Phi))$. Therefore, with $k = \frac{1}{2}(n + \operatorname{cross}(D(\Phi))$, Theorem 9 gives that $D(\Phi)$ can be untangled with at most k moves if and only if Φ is satisfiable. This gives the required **NP**-hardness via Lemma 3. (Note that $D(\Phi)$ and k can be obtained in polynomial time in the size of Φ.) ◀

The remainder of this section is devoted to the proof of Theorem 9.

5.1 Satisfiable implies small defect

In this subsection we show that given a satisfying assignment for Φ, $\operatorname{def}(D(\Phi)) \leq n$.

$\Phi = (x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee x_2 \vee \neg x_3) \wedge (\neg x_1 \vee \neg x_2 \vee x_4) \wedge (x_1 \vee \neg x_3 \vee \neg x_4)$

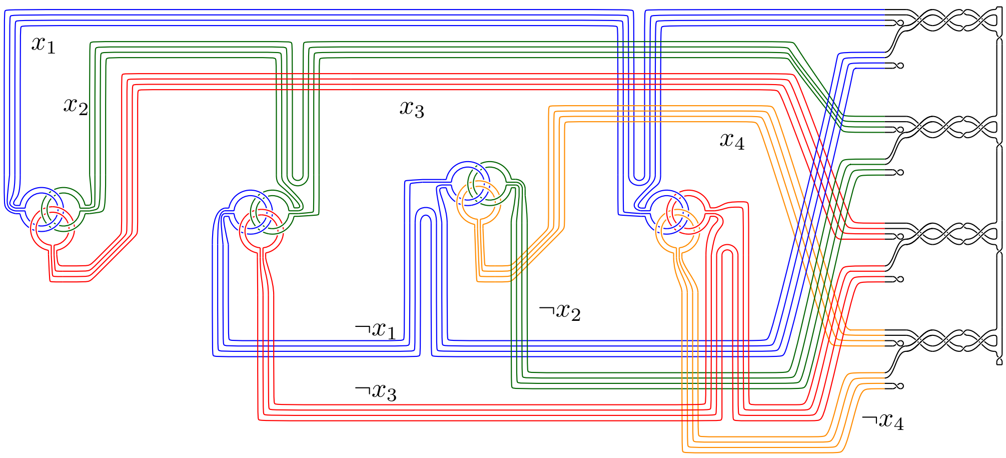

Figure 11 The final construction for the formula $\Phi = (x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee x_2 \vee \neg x_3) \wedge (\neg x_1 \vee \neg x_2 \vee x_4) \wedge (x_1 \vee \neg x_3 \vee \neg x_4)$. For simplicity of the picture, we do not visualize how the crossings are resolved in Step III. (Unfortunately, we cannot avoid tiny pictures of gadgets.)

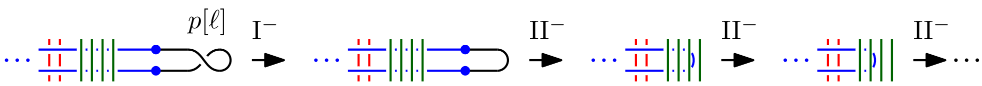

Figure 12 Initial simplifications.

For any literal $\ell = \text{TRUE}$, we remove the crossing $p[\ell]$ by a I^- move (see Figure 3). This uses exactly n I^- moves. Next we finish the untangling of the diagram by II^- moves only. Once this is done, we get an untangling with defect n by Lemma 7, completing the proof.

Thus it remains to finish the untangling with II^- moves only. For any $\ell = \text{TRUE}$, we shrink the ℓ tentacle by II^- moves (see Figure 12). By construction of $D(\Phi)$ we can continue shrinking the ℓ tentacle until we get a loop next to the $q[\ell]$ vertex; see Figure 13.

We continue the same process for every literal $\ell = \text{TRUE}$. Along the way, some of the arcs meeting $\gamma_1[\ell]$ and $\gamma_2[\ell]$ might have already been removed but it is still possible to simplify the ℓ tentacle as before. See Figure 14 for the result after shrinking all satisfied tentacles.

Because the assignment was satisfying, in each clause gadget at least one ring of the Borromean rings disappears. Consequently, if there are two remaining rings in some clause gadget, then they can be pulled apart from each other by II^- moves as in Figure 15.

After this step, for each $\ell = \text{TRUE}$, the $\gamma_1[\neg\ell]$ and $\gamma_2[\neg\ell]$ form "fingers" of four parallel curves that can be further simplified by II^- moves so that any crossings among different

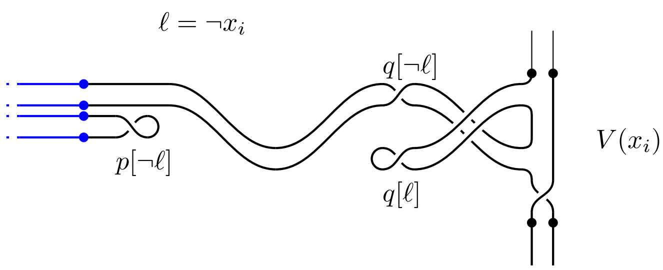

Figure 13 The ℓ tentacle was shrunk to a loop next to $q[\ell]$. In this example we have $\ell = \neg x_i$.

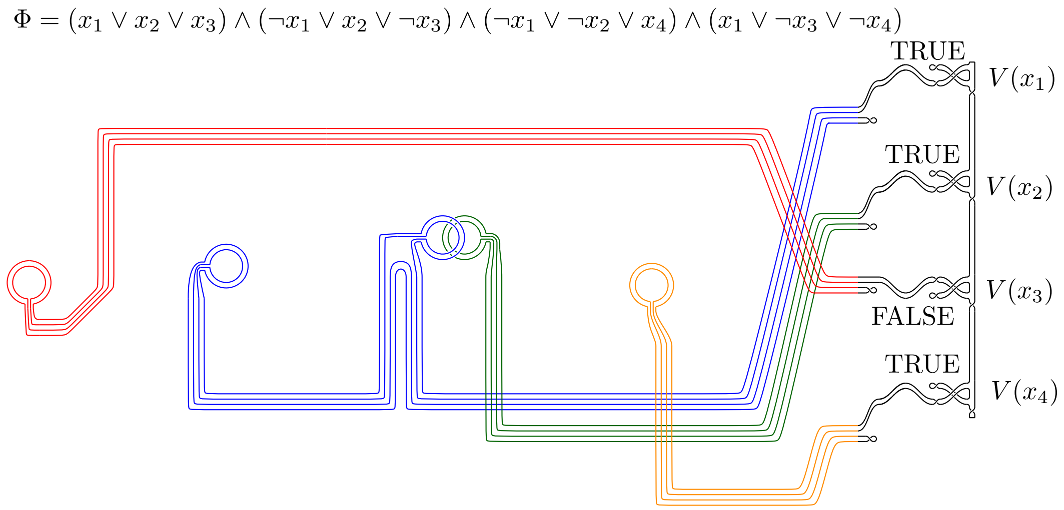

Figure 14 Simplified $D(\Phi)$ following a satisfying assignment $x_1 = x_2 = x_4 =$ TRUE and $x_3 =$ FALSE.

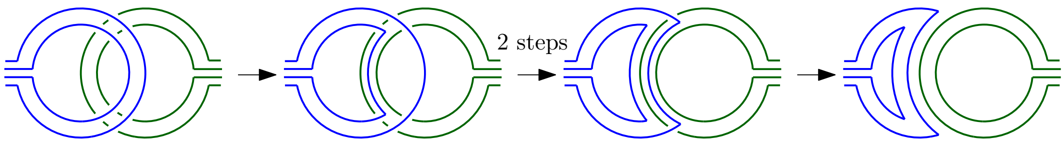

Figure 15 Untangling two rings in the clause gadget via II^- moves.

fingers are removed; see Figure 16. For each variable gadget we get one of the two pictures at Figure 17 left, and can be simplified to the picture on the right using II^- moves.

Finally, we recall how the variable gadgets are interconnected (compare the right picture at Figure 17 with Figure 10). Then it is easy to remove all remaining $2n$ crossings by II^- moves gradually from top to bottom. This finishes the proof of the "if" part of Theorem 9.

5.2 Small defect implies satisfiable

The purpose of this subsection is to sketch a proof of the "only if" part of the statement of Theorem 9. Recall that this means that we assume $\mathrm{def}(D(\Phi)) \leq n$ and we want to deduce that Φ is satisfiable. (Along the way we will actually also deduce that $\mathrm{def}(D(\Phi)) = n$.) In this subsection, we heavily use the terminology introduced in Section 4.

Let $\mathbf{D} = (D^0, \dots, D^k)$ be an untangling of D with $\mathrm{def}(\mathbf{D}) \leq n$. For a variable x let $R(x) = \{p[x], p[\neg x], q[x], q[\neg x], s_1(x), \dots, s_{12}(x)\}$ be the set of 16 out of the 17 self-crossings in the variable gadget $V(x)$ (we leave out $r(x)$) and let the *weight of* x, denoted by $w(x)$, be the sum of weights of the crossings in $R(x)$.

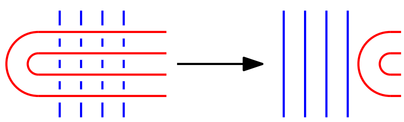

Figure 16 Simplifying $\gamma_1[\neg \ell]$ and $\gamma_2[\neg \ell]$ via II^- moves. First, we untangle the inner (horizontal) "finger" and then we untangle the outer (horizontal) "finger".

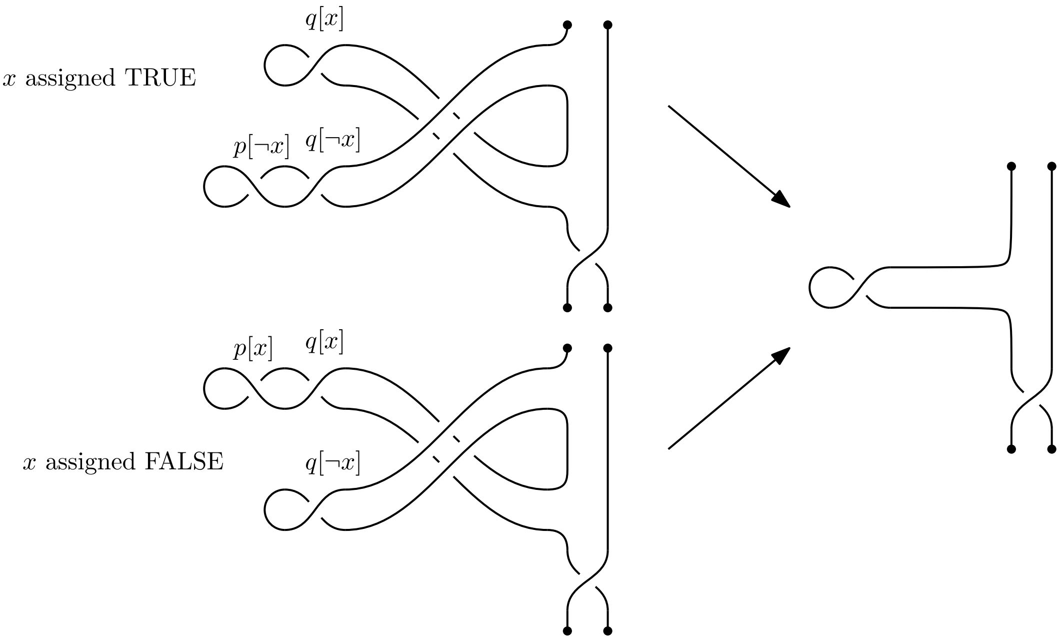

Figure 17 Results of the simplifications on the previous picture on the level of variable gadgets.

Now we provide two claims on the order of removals of the crossings. Both claims are proved using Lemma 8 via a careful case analysis, see Claims 10.1 and 10.2 of the full version [8]. The first claim analyzes the first economical II$^-$ move that removes some of the crossings in $R(x)$. For these claims, recall Figure 3.

▷ Claim 10. Let x be a variable with $w(x) \leq 1$. Let r be the first crossing in $R(x)$ which is removed by an economical II$^-$ move (we allow a draw). Then one of the following cases holds (note that in both cases $w(x) = 1$):

(i) $\{r, t(r)\} = \{s_1(x), s_2(x)\}$, $w(p[x]) = w(x) = 1$ and $p[x]$ is removed by a I$^-$ move prior to removing r and $t(r)$.

(ii) $\{r, t(r)\} = \{s_1(x), s_3(x)\}$, $w(p[\neg x]) = w(x) = 1$ and $p[\neg x]$ is removed by a I$^-$ move prior to removing r and $t(r)$.

Now, let us set $\ell := \ell(x) := x$ if the conclusion (i) of Claim 10 holds and $\ell := \ell(x) := \neg x$ if (ii) holds (assuming $w(x) \leq 1$). (We identify $\neg\neg x$ with x, that is, if $\ell = \neg x$, then $\neg \ell = x$.)

▷ Claim 11. If $w(x) \leq 1$, then $p[\neg \ell]$ and $q[\neg \ell]$ are twins. In addition, the preimage arcs α and β between $p[\neg \ell]$ and $q[\neg \ell]$ contain $\gamma_1[\neg \ell]$ and $\gamma_2[\neg \ell]$.

Now, we have acquired enough tools to finish the proof of the theorem. By Claim 10, we have $w(x) \geq 1$ for any variable x. By Lemma 7, we deduce

$$\operatorname{def}(\mathbf{D}) \geq \sum_x w(x) \geq n,$$

where the sum is over all variables. On the other hand, we assume $\operatorname{def}(\mathbf{D}) \leq n$. Therefore both inequalities above have to be equalities and in particular $w(x) = 1$ for any variable x. In particular, the assumptions of Claims 10 and 11 are satisfied for any variable x.

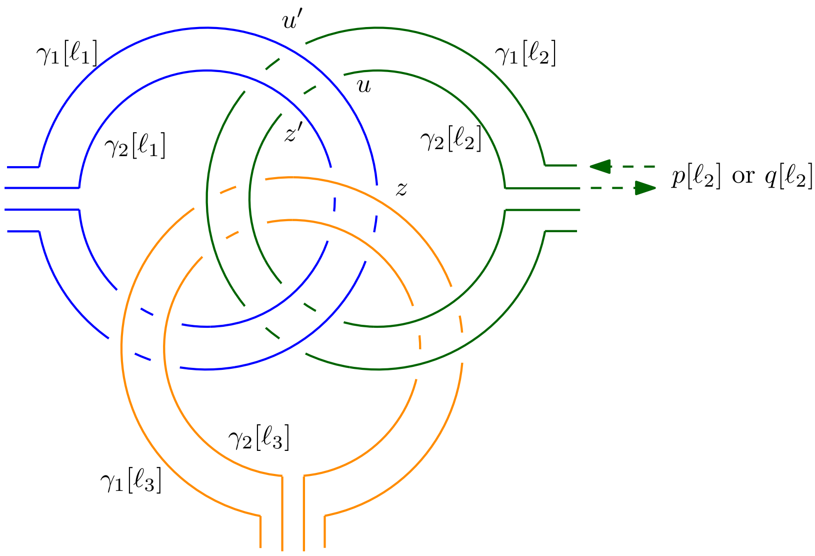

Figure 18 The clause gadget and some of the crossings of $R''(c)$.

Given a variable x, we assign x with TRUE if the conclusion (i) of Claim 10 holds (that is, if $x = \ell(x)$). Otherwise, if the conclusion (ii) of Claim 10 holds (i.e. $\neg x = \ell(x)$), we set x to FALSE. It remains to prove that we get a satisfying assignment this way.

For contradiction, suppose there is a clause $c = (\ell_1 \vee \ell_2 \vee \ell_3)$ which is not satisfied with this assignment. Let x_i be the variable of ℓ_i, that is, $\ell_i = x_i$ or $\ell_i = \neg x_i$. The fact that c is not satisfied with the assignment above translates as $\ell(x_i) = \neg \ell_i$ for any $i \in \{1, 2, 3\}$.

By Claim 11, we get that $p[\ell_i]$ and $q[\ell_i]$ are twins for any $i \in \{1, 2, 3\}$. Let $R''(c)$ be the set of crossings in the clause gadget of c union the sets $\{p[\ell_i], q[\ell_i]\}$ for $i \in \{1, 2, 3\}$. All the crossings in $R''(c)$ have weight 0 and they have to be removed by economical II^- moves as all defect is realized on points $p[\ell(x)]$ (but $p[\ell_i] = p[\neg \ell(x_i)]$ are not among these points).

Let r be the first removed crossing among the crossings in $R''(c)$. Our aim is to rule out all options in which a crossing of $R''(c)$ can be r. This way, we will obtain the required contradiction. First, we observe that r cannot be any of $p[\ell_i]$ or $q[\ell_i]$ for $i \in \{1, 2, 3\}$. This follows from Claim 11 as the arcs $\gamma_1[\ell_i]$ and $\gamma_2[\ell_i]$ contain some crossings in $R''(c)$.

Next we apply Lemma 8 (for 0-close neighbors) with $R = R''(c)$. For sake of example, in this sketch we rule out the (interesting) option $r = u$, where u is the vertex on Figure 18. Let α and β be the arcs between r and $t(r)$ from the definition of 0-close neighbors (α and β exist by Lemma 8). In particular, neither α nor β has a crossing from $R''(c)$ in its interior. The only option for α is to emanate to the left reaching the crossing u' as emanating to the right reaches a crossing $z \in R''(c)$ as an underpass. In particular, $t(u) = u'$. Consequently, β has to emanate to the right since emanating to the left would reach z'. However, before β reaches u', it has to pass through $p[\ell_2]$ or $q[\ell_2]$ in $R''(c)$ which rules out this option.

6 Intermediate invariants

In this section we describe the family of link invariants from the statement of Theorem 2. The material presented here is standard and details can be found in various textbooks; since every piecewise linear knot can be smoothed in a unique way, we assume that the knots discussed are smooth. In Sections 6 and 7 we work in the smooth category. We first define:

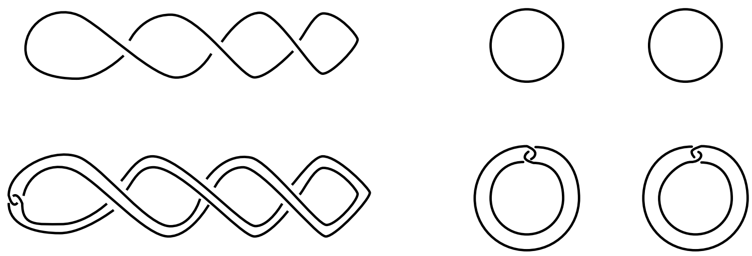

Figure 19 Some Whitehead doubles.

▶ **Definition 12.** *Let L be a link in the 3-sphere. We now give a list of the invariants that we will be using; for a detailed discussion see, for example, [23].*

1. *A* smooth slice surface *for L is an orientable surface with no closed components, properly and smoothly embedded in the 4-ball, whose boundary is L.*
2. *The* 4-ball Euler characteristic *of L, denoted $\chi_4(L)$, is the largest integer so that L bounds a smooth slice surface of Euler characteristic $\chi_4(L)$.*
3. *A link is called* smoothly slice *if it bounds a slice surface that consists entirely of disks; equivalently, the $\chi_4(L)$ equals the number of components of L.*
4. *The* unlinking number*, denoted $u(L)$, is the smallest nonnegative integer so that L admits some diagram D so that after $u(L)$ crossing changes on D a trivial link is obtained.*
5. *By transversality, every link L bounds smoothly immersed disks in B^4 with finitely many double points. The* 4-dimensional clasp number *(sometimes called the* 4-ball crossing number*) of L, denoted $c_s(L)$, is the minimal number of double points for such disks.*

Finally, we define intermediate invariants, whose existence follows from Theorem 2 of [25].

▶ **Definition 13** (intermediate invariant)**.** *A real valued link invariant $i(L)$ is called an* intermediate invariant *if $u(L) \geq i(L) \geq c_s(L)$.*

Many invariants are known to be intermediate (see, for example, [25]). We list a few (and prove that they are intermediate) in Lemma 13 of the full version [8]. Shibuya [25] proved:

▶ **Lemma 14.** *Let L be a link with μ components. Then $\chi_4(L) \geq \mu - 2c_s(L)$.*

7 Unlinking, 4-ball Euler characteristic, and intermediate invariants

In this section we will show that the link invariants defined in the previous section are **NP**-hard. Our reduction relies on the use of Whitehead doubles, which we now define:

▶ **Definition 15** (Whitehead Double)**.** *Let L be a link. A* Whitehead double *of L is a link obtained by taking two parallel copies of each component of L and joining them together with a clasp (see Figure 19). A Whitehead double is called* positive *if the crossings at the clasp are positive. If the linking number of the two copies of each component is zero the Whitehead double is called* untwisted*. It is easy to see that the untwisted positive Whitehead double is uniquely determined by L.*

One last piece of background we will need is a result of Levine [22, Theorem 1.1]:

▶ **Lemma 16.** *The untwisted positive Whitehead double of the Hopf link, and that of the Borromean rings, are not smoothly slice.*

We are now ready to describe our construction:

The construction of L_Φ^{WH}. Given a 3-SAT instance Φ, recall the link L_Φ from Section 3, and let L_Φ^{WH} be its positive untwisted Whitehead double. There is a natural bijection between components before and after taking a Whitehead double; let $\kappa_{x_i}^{\text{WH}}$ denote the component corresponding to κ_{x_i} and let $\kappa_{\neg x_i}^{\text{WH}}$ denote the component corresponding to $\kappa_{\neg x_i}$

The goal of this section is to prove:

▶ **Theorem 17.** *Given a 3-SAT instance Φ with n variables, let L_Φ^{WH} be the link constructed above. Then the following are equivalent, where here i is any intermediate invariant:*

1. *Φ is satisfiable.*
2. *$u(L_\Phi^{WH}) = n$.*
3. *$i(L_\Phi^{WH}) = n$.*
4. *$c_s(L_\Phi^{WH}) = n$.*
5. *$\chi_4(L_\Phi^{WH}) = 0$.*
6. *L_Φ^{WH} admits a smoothly slice sublink with n components.*

Theorem 2(b) – (d) directly follows from Theorem 17.

Proof. The proof of this Theorem is split in the following steps:

(a) Φ is satisfiable implies that $u(L_\Phi^{\text{WH}}) \leq n$.

(b) $c_s(L_\Phi^{\text{WH}}) \leq i(L_\Phi^{\text{WH}}) \leq u(L_\Phi^{\text{WH}})$.

(c) $\chi_4(L_\Phi^{\text{WH}}) \geq 2n - 2c_s(L_\Phi^{\text{WH}})$.

(d) If $\chi_4(L_\Phi^{\text{WH}}) \geq 0$ then the following two conditions hold:

  (d.I) $\chi_4(L_\Phi^{\text{WH}}) = 0$

  (d.II) L_Φ^{WH} admits a smoothly slice sublink with n components.

(e) If L_Φ^{WH} admits a smoothly slice sublink with n components, then Φ is satisfiable.

We first show how (a) – (e) prove Theorem 17. Assume first that Φ is satisfiable. Then by (a) and (b) we have that $c_s(L_\Phi^{\text{WH}}) \leq n$, and by (c) we have that $\chi_4(L_\Phi^{\text{WH}}) \geq 0$. Then (d.I) shows that $\chi_4(L_\Phi^{\text{WH}}) = 0$. Working our way back, we see that $c_s(L_\Phi^{\text{WH}}) = i(L_\Phi^{\text{WH}}) = u(L_\Phi^{\text{WH}}) = n$, establishing (1) $\Rightarrow$ (2) $\Rightarrow$ (3) $\Rightarrow$ (4) $\Rightarrow$ (5). In addition, (d.II) shows directly that (5) $\Rightarrow$ (6). Finally, (e) establishes (6) $\Rightarrow$ (1).

We complete the proof of Theorem 17 by establishing (a) – (e):

(a) Suppose we have a satisfying assignment for Φ (for this implication, *cf.* the proof of Theorem 5). By a single crossing change we resolve the clasp of every component that correspond to a satisfied literal, that is, if x_i = TRUE we change one of the crossings of the clasp of $\kappa_{x_i}^{\text{WH}}$ and if x_i = FALSE we do it for $\kappa_{\neg x_i}^{\text{WH}}$; as a result, the components corresponding to satisfied literals now form unlinks that are not linked with the remaining components, and we can isotope them away. Since the assignment is satisfying, from each copy of the Borromean rings at least one ring is removed, and the remaining components retract into the first n disks that contained the Hopf links. In each disk we have an untwisted Whitehead double of the unknot which is itself an unknot. Thus we see that the unlink on $2n$ component is obtained, showing that $u(L_\Phi^{\text{WH}}) \leq n$.

(b) By definition of intermediate invariant we have that $c_s \leq i \leq u$.

(c) This is Lemma 14.

(d) This step is the crux of the proof and relies on the computation of the *signature* of our link; see step (d) in the proof of Theorem 20 of the full version [8].

(e) Finally, assume that L_Φ^{WH} admits an n component smoothly slice sublink L_{slice}. By Lemma 16 we have that the positive untwisted Whitehead double of the Hopf link is not a sublink of L_{slice} and therefore for each i exactly one of $\kappa_{x_i}^{\text{WH}}$ and $\kappa_{\neg x_i}^{\text{WH}}$ is in L_{slice}. If $\kappa_{x_i}^{\text{WH}}$ is in L_{slice} we set $x_i = \text{FALSE}$ and if $\kappa_{\neg x_i}^{\text{WH}}$ is in L_{slice} we set $x_i = \text{TRUE}$. Using Lemma 16 again we see that L_{slice} does not admit the positive untwisted Whitehead double of the Borromean rings as a sublink. Therefore, from every set of Borromean rings, at least one component does *not* belong to L_{slice}; therefore the assignment is satisfying. ◀

References

1. Ian Agol, Joel Hass, and William Thurston. The computational complexity of knot genus and spanning area. *Transactions of the American Mathematical Society*, 358:3821–3850, 2006. doi:10.1090/S0002-9947-05-03919-X.
2. Francesca Aicardi. Tree-like curves. *Singularities and bifurcations*, 21:1–31, 1994.
3. Vladimir I Arnold. Plane curves, their invariants, perestroikas and classifications. *Advances in Soviet Mathematics*, 21:33–91, 1994.
4. Vladimir Igorevich Arnold. *Topological invariants of plane curves and caustics*, volume 5. American Mathematical Soc., 1994.
5. Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, Cambridge, 2009. doi:10.1017/CBO9780511804090.
6. Hsien-Chih Chang and Jeff Erickson. Untangling planar curves. *Discrete & Computational Geometry*, 58(4):889–920, 2017. doi:10.1007/s00454-017-9907-6.
7. Arnaud de Mesmay, Yo'av Rieck, Eric Sedgwick, and Martin Tancer. Embeddability in $\mathbb{R}^3$ is NP-hard. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1316–1329. Society for Industrial and Applied Mathematics, 2018. Full version on arXiv:1604.00290. doi:10.1137/1.9781611975031.86.
8. Arnaud de Mesmay, Yo'av Rieck, Eric Sedgwick, and Martin Tancer. The unbearable hardness of unknotting. arXiv:1810.03502 (for numbering of claims, we refer to the version v1), 2018.
9. Michael R. Garey and David S. Johnson. *Computers and intractability*. W. H. Freeman and Co., San Francisco, Calif., 1979. A guide to the theory of NP-completeness, A Series of Books in the Mathematical Sciences.
10. Wolfgang Haken. Theorie der Normalflächen. *Acta Mathematica*, 105(3-4):245–375, 1961.
11. Joel Hass and Jeffrey Lagarias. The number of Reidemeister moves needed for unknotting. *Journal of the American Mathematical Society*, 14(2):399–428, 2001. doi:10.1090/S0894-0347-01-00358-7.
12. Joel Hass, Jeffrey C. Lagarias, and Nicholas Pippenger. The computational complexity of knot and link problems. *Journal of the ACM (JACM)*, 46(2):185–211, 1999. doi:10.1145/301970.301971.
13. Joel Hass and Tahl Nowik. Unknot diagrams requiring a quadratic number of Reidemeister moves to untangle. *Discrete & Computational Geometry*, 44(1):91–95, 2010. doi:10.1007/s00454-009-9156-4.
14. Louis H Kauffman and Sofia Lambropoulou. Hard unknots and collapsing tangles. In *Introductory Lectures on Knot Theory*, 2014. doi:10.1142/9789814313001_0009.
15. Dale Koenig and Anastasiia Tsvietkova. NP-hard problems naturally arising in knot theory, 2018. arXiv:1809.10334.
16. Greg Kuperberg. Knottedness is in **NP**, modulo GRH. *Adv. Math.*, 256:493–506, 2014. doi:10.1016/j.aim.2014.01.007.
17. Greg Kuperberg and Eric Samperton. Coloring invariants of knots and links are often intractable. Manuscript.
18. Marc Lackenby. A polynomial upper bound on Reidemeister moves. *Ann. Math. (2)*, 182(2):491–564, 2015. doi:10.4007/annals.2015.182.2.3.

19 Marc Lackenby. The efficient certification of knottedness and Thurston norm, 2016. `arXiv:1604.00290`.

20 Marc Lackenby. Elementary knot theory. In *Lectures on Geometry (Clay Lecture Notes)*. Oxford University Press, 2017.

21 Marc Lackenby. Some conditionally hard problems on links and 3-manifolds. *Discrete & Computational Geometry*, 58(3):580–595, 2017. `doi:10.1007/s00454-017-9905-8`.

22 Adam Simon Levine. Slicing mixed Bing–Whitehead doubles. *Journal of Topology*, 5(3):713–726, 2012. `doi:10.1112/jtopol/jts019`.

23 Dale Rolfsen. *Knots and links*, volume 7 of *Mathematics Lecture Series*. Publish or Perish, Inc., Houston, TX, 1990. Corrected reprint of the 1976 original.

24 Eric Samperton. Computational Complexity of Enumerative 3-Manifold Invariants. *arXiv preprint*, 2018. `arXiv:1805.09275`.

25 Tetsuo Shibuya. Some relations among various numerical invariants for links. *Osaka J. Math.*, 11:313–322, 1974. URL: `http://0-projecteuclid.org.library.uark.edu/euclid.ojm/1200757391`.