

On the number of unknot diagrams

CAROLINA MEDINA¹, JORGE L. RAMÍREZ-ALFONSÍN^{2,3}, AND GELASIO SALAZAR^{1,3}

¹*Instituto de Física, UASLP. San Luis Potosí, Mexico, 78000.*

²*Institut Montpellierain Alexander Grothendieck, Université de Montpellier. Place Eugène Bataillon, 34095 Montpellier, France.*

³*Unité Mixte Internationale CNRS-CONACYT-UNAM “Laboratoire Solomon Lefschetz”. Cuernavaca, Mexico.*

October 17, 2017

Abstract

Let D be a knot diagram, and let \mathcal{D} denote the set of diagrams that can be obtained from D by crossing exchanges. If D has n crossings, then \mathcal{D} consists of 2^n diagrams. A folklore argument shows that at least one of these 2^n diagrams is unknot, from which it follows that every diagram has finite unknotting number. It is easy to see that this argument can be used to show that actually \mathcal{D} has more than one unknot diagram, but it cannot yield more than $4n$ unknot diagrams. We improve this linear bound to a superpolynomial bound, by showing that at least $2^{\sqrt[3]{n}}$ of the diagrams in \mathcal{D} are unknot. We also show that either all the diagrams in \mathcal{D} are unknot, or there is a diagram in \mathcal{D} that is a diagram of the trefoil knot.

1 Introduction

We follow standard knot theory terminology as in [1]. All knots under consideration are tame, that is, they have regular diagrams. All diagrams under consideration are regular.

1.1 Our main result

A staple in every elementary knot theory course is that, given any diagram D , it is always possible to turn D into a diagram of the unknot (that is, into an *unknot diagram*) by performing some crossing exchanges on D (see Figure 1). The minimum number of such crossing exchanges that turn D into an unknot diagram is a well-studied parameter, the *unknotting number* of D .

We review in Section 2 the well-known argument that turns any diagram D into an unknot diagram, by performing crossing exchanges. As we shall remark, this argument actually produces more than one unknot diagram from D , but it gives a relatively small number: if D has n crossings, this idea yields at most $4n$ distinct unknot diagrams that can be obtained from D by crossing exchanges. Our main result is that actually there is always a superpolynomial number of such unknot diagrams.

Theorem 1. *Let D be a knot diagram with n crossings. Then there are at least $2^{\sqrt[3]{n}}$ distinct unknot diagrams that can be obtained by performing crossing exchanges on D .*

As far as we know, no non-trivial general lower bound along the lines of Theorem 1 has been previously reported. There is, however, a substantial amount of work related to the complementary problem: out of the 2^n diagrams that can be obtained from D by crossing exchanges, how many are knotted? This is related to the well-studied Frisch-Wasserman-Delbruck conjecture (see Section 2), which in its general form can be paraphrased as follows: if we take a random knot diagram with n crossings, then the probability that it is knotted tends to 1 as n goes to infinity.

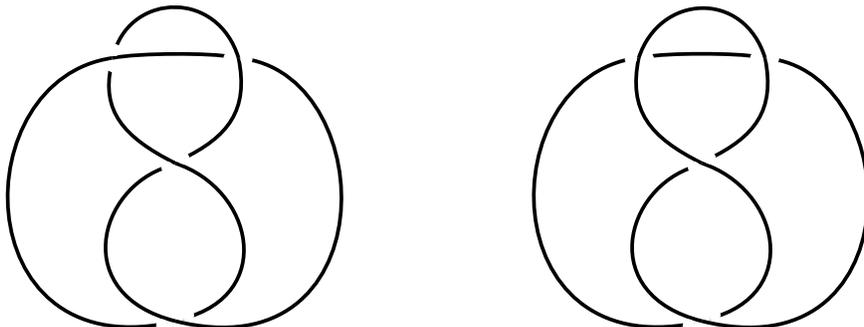


Figure 1: On the left hand side we have a diagram of the figure-eight knot. By performing a crossing exchange on the top left crossing, we obtain the unknot diagram on the right hand side.

For our arguments, it will be convenient to formulate Theorem 1 in terms of shadows. We recall that the *shadow* S of a diagram D is the 4-regular plane graph obtained by transforming each crossing of D into a vertex.

Conversely, we may regard a diagram as a shadow for which we provide the over/under-crossing information at each vertex. That is, for each vertex it is prescribed which strand goes locally above (or, equivalently, is the *overpass*) and which strand goes locally below (or, equivalently, is the *underpass*). We call this information the *prescription* at the vertex.

At each vertex we have two possible prescriptions (corresponding to which of the strands is an overpass at the corresponding crossing). The collection of prescriptions for all the vertices yields an *assignment* on S , resulting in a diagram D . We say that S is the shadow of D , and that D is a diagram of S . If S has n vertices, then there are 2^n possible assignments on S , and so there are 2^n distinct diagrams of S . These 2^n distinct diagrams are the diagrams that can be obtained by performing crossing exchanges on D .

In view of this discussion, it follows that Theorem 1 can be alternatively written as follows.

Theorem 2 (Equivalent to Theorem 1). *Let S be a shadow with n vertices. Then S has at least $2^{\sqrt[3]{n}}$ distinct unknot diagrams.*

We use $\mathcal{U}(S)$ to denote the set of unknot diagrams of S . Thus, with this notation, Theorem 2 claims that every shadow S satisfies $|\mathcal{U}(S)| \geq 2^{\sqrt[3]{n}}$.

Obviously, the exact number of unknot diagrams depends on the shadow under consideration. It is easy to show that out of the 2^4 diagrams obtained from the shadow of the diagrams in Figure 1, exactly 12 are unknot. There are also extreme examples such as the ones illustrated in Figure 2: all 2^4 diagrams of each of these shadows are unknot. It is straightforward to generalize these examples, showing that for every positive integer n there is a shadow on n vertices, all of whose diagrams are unknot.

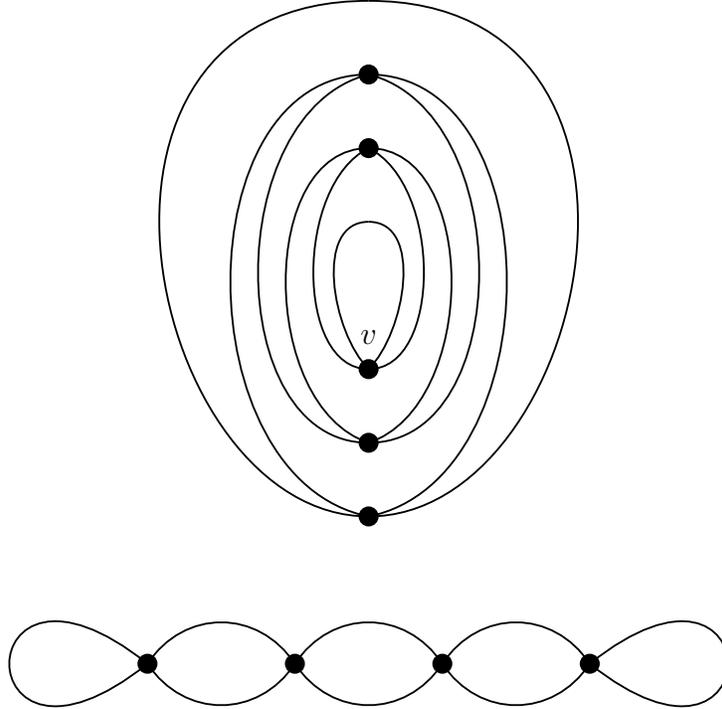


Figure 2: All the diagrams of these shadows are unknot.

1.2 Shadows that yield the trefoil knot

We also investigate which knots (other than the unknot) can be guaranteed to be obtainable from a given shadow. That is, given a knot K and a shadow S , is there an assignment on S that gives a diagram of K ?

The first natural step is the trefoil knot. If D is a diagram whose corresponding knot is equivalent to the trefoil knot, then we say that D is a *trefoil diagram*.

For this particular knot we can give a complete answer. First let us note that there exist arbitrarily large shadows that do not have trefoil diagrams: it suffices to generalize the examples in Figure 2 to an arbitrary number of vertices. What we will show is that such shadows (that do not have trefoil diagrams) can be easily characterized.

A glaring feature of the shadows in Figure 2 is that each of their vertices is a cut-vertex. We use Tutte's notion of a cut-vertex [29], which we review in Section 9. We show that the property that every vertex is a cut-vertex completely characterizes which shadows have only unknot diagrams, and that every shadow that does not have this feature has a trefoil diagram:

Theorem 3. *Let S be a shadow. Then every diagram of S is unknot if and only if every vertex of S is a cut-vertex. Moreover, if not every vertex of S is a cut-vertex, then S has a trefoil diagram.*

It is natural to ask how far we can get along the lines of Theorem 3. The next obvious step is to investigate the figure-eight knot, which is the only knot with crossing number four. One might think that for the figure-eight knot, a result analogous to Theorem 3 could hold, at least for all sufficiently large non-simple shadows.

Now it is easy to construct an arbitrarily large non-simple shadow, none of whose assignments gives a diagram of the figure-eight knot. Indeed, it suffices to artificially grow the trefoil diagram by gluing to it the shadow at the bottom of Figure 2, as shown in Figure 3.

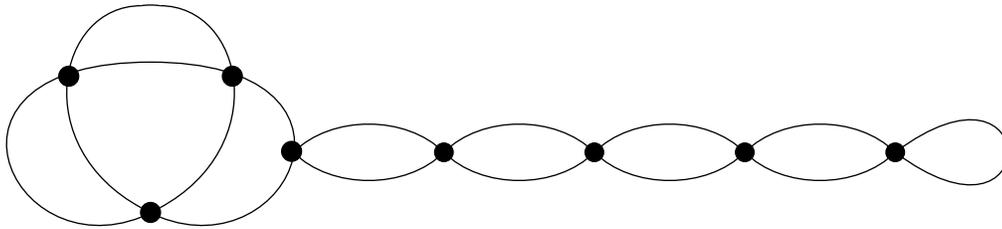


Figure 3: Every diagram of this shadow is either a trefoil diagram or an unknot diagram, and this can be easily generalized to obtain arbitrarily large shadows with the same property.

Thus, in order to meaningfully explore a possible extension of Theorem 3 to shadows other than the trefoil (say, to the figure-eight knot) one must exclude from consideration shadows that have cut vertices; that is, we must focus our attention on 2-connected shadows.

As it happens, no analogue of Theorem 3 holds for the figure-eight knot, even if the attention is restricted to 2-connected shadows, no matter how large. Moreover, no such analogue exists for any knot with positive even crossing number:

Observation 4. *Let K be any knot with positive even crossing number, and let n be a positive integer. Then there exists a 2-connected shadow S with more than n vertices such that, for every diagram D of S , the knot corresponding to D is not equivalent to K .*

1.3 Outline of the proof of Theorem 2 and structure of this paper

In the next section we overview some previous results in the literature related to Theorem 2, and in Section 3 we review some elementary facts on shadows that we use extensively in this paper.

Our proof of Theorem 2 is of an inductive nature. In a nutshell the idea is that, given a shadow S , we obtain from S a smaller shadow S' that is (loosely speaking) contained in S , and then show that every unknot diagram of S' can be extended to at least two unknot diagrams of S . Applying a similar procedure to S' , by an inductive reasoning we find the claimed lower bound for $\mathcal{U}(S)$.

For an illustration of one of the techniques we use, we refer the reader to Figure 7. If we remove from S the edges of the thick cycle C and suppress all resulting degree 2 vertices, we obtain the shadow S' shown below S . We will use the notation $S' = S//C$ to indicate that S' is obtained from S in this way. To the right of S' we show an unknot diagram D' of S' , and above D' we show how to extend D' to four unknot diagrams of S . The idea here is simply that it suffices to have the strand corresponding to C lie entirely above or entirely below the rest of the diagram. This property implies that D' and D are equivalent; since D' is unknot, then these four diagrams are unknot.

We note that this technique cannot be applied using an arbitrary cycle C . We can use it only if C is a *straight-ahead* cycle, that is, as we traverse C (starting from some of its vertices), then, as we reach an intermediate vertex v , we must continue traversing C using the edge opposite to the one from which we arrived to v . (Straight-ahead cycles are called loops in [14]).

In order to simplify our arguments, we will use the notion of a *macro move* in a diagram (following the terminology used by Kauffman in [21]). An instance of a macro move is precisely when we have a strand all of whose crossings lie “above” the rest of the diagram; such a strand is an *overstrand*. Similarly, if all the crossings of the strand lie “below” the rest of the diagram, then the strand is an *understrand*. An overstrand (respectively, understrand) can be replaced with another overstrand (respectively, understrand) with the same endpoints, if these endpoints are distinct; the diagram thus obtained is equivalent to the original diagram. If an overstrand or understrand starts and ends at the same crossing point, then we can collapse the strand to this point, and the diagram thus obtained is equivalent to the original diagram. This is precisely the operation performed on the thick strands in diagrams D_1, D_2, D_3 , and D_4 in Figure 7. We will review macro moves in Section 4.

Back to the discussion on the proof of Theorem 2, we illustrate the second main technique we use in Figure 12. In the shadow S we identify two parallel edges with common endpoints u and v ; these two edges and their endpoints form a *digon*. We obtain S' by splitting the vertices u and v as shown. Chang and Erickson call this type of operation a $2 \rightarrow 0$ homotopy move, and this suggests that S' has two fewer vertices than S . To the right of S' we have an unknot diagram D' of S' . This diagram can be extended to the two illustrated diagrams D_1 and D_2 of S . In each of D_1 and D_2 we can apply a Reidemeister move of Type II on the strands corresponding to the identified digon, obtaining D' . Since D' is unknot, it follows that both D_1 and D_2 are unknot.

In Section 5 we formally explain and develop this first technique we have outlined. The limitation of the first technique is that, in order to successfully apply it recursively, it must be used on straight-ahead cycles that have a relatively small number of vertices: if the cycle whose edges we remove has a large number of vertices, then the resulting shadow S' will have too few vertices, and so the inductive argument will yield a poor lower bound for $|\mathcal{U}(S)|$. The problem is that we cannot guarantee that there always exist (relatively) small straight-ahead cycles. It is easy to give examples of arbitrarily large shadows for which there are only two cycles on which this technique can be used, and getting rid of one of these cycles trivializes the whole shadow.

For such cases in which the first technique does not suffice, we have at hand the second technique. Now the obvious limitation of this second technique is that there exist shadows that do not have any digons. To handle this possibility we need a variant of this second technique, which can be applied to a pair of internally disjoint paths with common endpoints. Such a pair of paths has a natural underlying digon-like structure, and most of the work needed to prove Theorem 2 involves the extension of the second technique to take advantage of such digon-like structures. Part of this work involves a detailed understanding on the existence and structure of digons in a system of two simple closed curves in the plane. This is carried out in Section 6. With this information at our disposal, we then formally lay out the second technique in Section 7.

The final step, given in Section 8, consists then on showing how these two techniques can be combined to yield Theorem 2.

The proof of Theorem 3 is in Section 9. In a nutshell, to prove Theorem 3 we investigate the structure of a shadow not all of whose vertices are cut vertices, and then explicitly construct a trefoil diagram for any such shadow. Also in this section we describe the infinite family of shadows that proves Observation 4. Finally, in Section 10 we present some concluding remarks and open questions.

2 Related work

Let us start by recalling the well-known argument showing that if S is a shadow, then there is an assignment on S that gives an unknot diagram. Let S be a shadow, and let p be a point in the middle of an edge of S . Starting at p , we traverse the curve underlying S , following any of the two possible directions (see Figure 4(a)). We obtain a diagram from S by following the rule that each time we arrive to a vertex for the first time, then we assign a crossing to this vertex so that the strand we are currently traversing is the overstrand. If we apply this procedure to the shadow in Figure 4(a), we get the unknot diagram shown in Figure 4(b). It is easy to convince oneself that a diagram produced in this manner is unknot; a formal proof can be found in [1, Section 3.1].

An unknot diagram obtained in this way is called a *standard unknot diagram* (see [20]), and also a *descending diagram* (see [22] and [24]). To understand the motivation behind the terminology “descending”, let us consider a knot that gets projected to such a diagram. Suppose that the z -axis is the one orthogonal to the plane where the diagram is drawn. Then the knot can be realized (in 3-dimensional space) so that, as we traverse the knot starting at the point p' that projects to p , the z -coordinates are descending until we get to a point q' whose projection is very close to p , and finally we join q' to p' with a z -ascending strand.

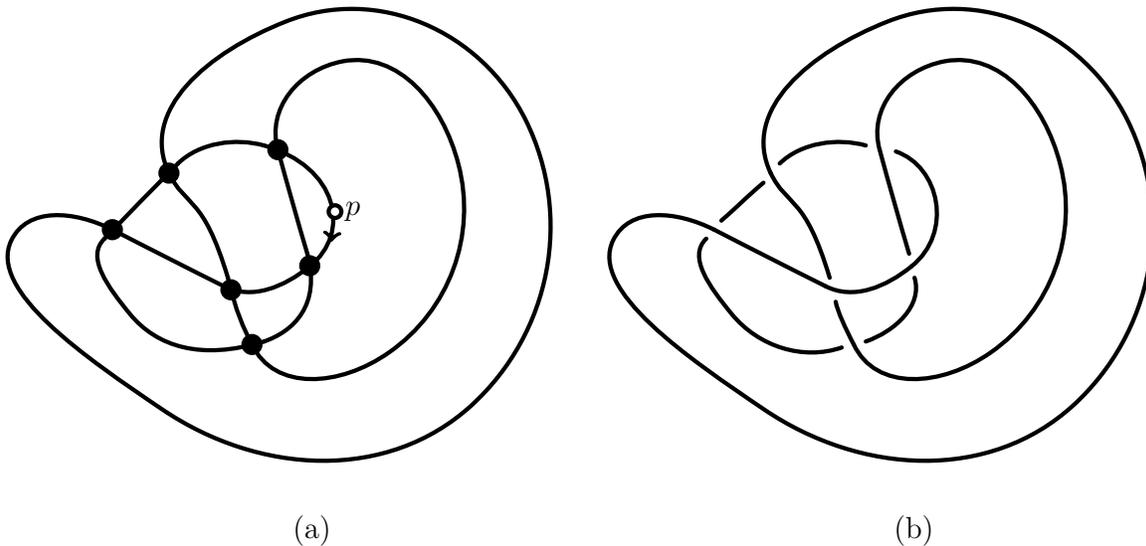


Figure 4: Creating an unknot diagram starting from any given shadow.

As we mentioned in Section 1, for any given shadow S with n vertices, the number of unknot diagrams of S that can be produced with this construction is at most $4n$. Indeed, there are $2n$ ways to choose p (since S has $2n$ edges), and for each of these $2n$ choices there are two possible directions to follow.

There seem to be very few results in the literature involving the number of unknot diagrams of an arbitrary shadow. Recently, Cantarella et al. [5] carried out a complete investigation of the knot types that arise in all shadows with 10 or fewer crossings. Their results show that, among all the diagrams associated to this huge collection of shadows, roughly 78% are unknot diagrams. Cantarella et al. observed that this large proportion is explained by the existence of “tree-like”

shadows, such as the ones depicted in Figure 2. These are shadows in which every vertex is a cut-vertex, and so an easy induction shows that all of their diagrams are unknot.

There is a collection of quite interesting, deep results in the literature related to the complementary aspect of Theorem 2: out of the 2^n diagrams that have S as their shadow, how many are *not* unknot diagrams?

According to Sumners and Whittington [27], this was first asked in the context of long linear polymer chains, independently by Frisch and Wasserman [16] and Delbruck [9]. In our context, the Frisch-Wasserman-Delbruck Conjecture may be roughly paraphrased as follows: if $p(n)$ is the probability that a random knot diagram with n crossings is knotted, then $p(n) \rightarrow 1$ as $n \rightarrow \infty$.

The first issue to settle in an investigation of the Frisch-Wasserman-Delbruck Conjecture is: what is a random knot diagram? In [27], Sumners and Whittington investigated (and settled in the affirmative) this conjecture in the model of self-avoiding walks on the three-dimensional simple cubic lattice (see [25] for a closely related result). The conjecture has also been settled in other models of space curves, such as self-avoiding Gaussian polygons and self-avoiding equilateral polygons [3, 4, 11, 12, 18].

The problem of proposing suitable models of random knot diagrams is of interest by itself. The associated difficulties are explained and discussed in [10], where two different such models are presented and investigated. We also refer the reader to the preliminary report by Dunfield et al. in [13]. It is also worth mentioning the very recent work of Even-Zohar et al. [15], where several rigorous results are established on the distributions of knot and link invariants for the Petaluma model, which is based on the representation of knots and links as petal diagrams [2]. In a very interesting recent development, Chapman [6, 7] proved the Frisch-Wasserman-Delbruck conjecture under a very general and natural model.

From the work of Chapman we know that if we take a random diagram on n vertices, then the probability that it is an unknot diagram decreases exponentially with n . Our work in this paper focuses on the abundance of unknot diagrams associated to an arbitrary (that is, not to a random) shadow.

3 Basic facts on shadows

A shadow of a diagram with at least one crossing is a 4-regular plane graph. For our upcoming arguments it will be convenient to admit a vertex-less simple closed curve as a shadow. This *trivial* shadow is the shadow of a *trivial* diagram, that is, a diagram with no crossings.

Thus every shadow is a 4-regular plane graph, but not every 4-regular plane graph is a shadow of a knot. A 4-regular graph is a shadow if and only if it has a straight-ahead Eulerian closed walk. We recall that a *straight-ahead* walk in a 4-regular graph is a walk in which every time we reach a vertex v in the walk, the next edge in the walk is the opposite edge to the one from which we arrived to v . A walk is *closed* if its final vertex is the same as its initial vertex, and it is *Eulerian* if each edge of the graph gets traversed exactly once in the graph.

We also recall that if C is a connected subgraph of a graph G , and every vertex of C has degree two (in C), then C is a *cycle* of G . For our purposes, a vertex-less graph (which, we recall, consists of a single edge, homeomorphic to a simple closed curve) will be considered a cycle.

If C is a cycle of a shadow S , and there exists a straight-ahead walk that traverses each edge of C exactly once, then C is a *straight-ahead* cycle of S . The initial (and final) vertex of this walk is a *root* of C . For an illustration of a straight-ahead cycle in a shadow we refer the reader to the top left

of Figure 7. The trivial shadow has no vertices, but by convention we regard it as a straight-ahead cycle. Finally we note that it is easy to see that a straight-ahead cycle in a nontrivial shadow has a unique root.

In our arguments we make extensive use of the following properties of shadows. The first one is totally straightforward.

Fact 5. *Let v be a vertex of a shadow S . Then there is an Eulerian straight-ahead closed walk W that starts and ends at v . The walk W gets naturally decomposed into two edge-disjoint straight-ahead closed walks W_1, W_2 , both of which have v as their startvertex and endvertex.*

Fact 6. *Every nontrivial shadow has at least two distinct (necessarily edge-disjoint) straight-ahead cycles.*

Proof. Let v be a vertex of a shadow S , and let W_1, W_2 be the straight-ahead closed walks (from Fact 5) that start and end at v . We show that as we traverse W_1 , we must encounter a straight-ahead cycle. Since the argument is symmetric on W_1 and W_2 , and these walks are edge-disjoint, this proves the fact.

We prove this by induction on the number of vertices of the walk. If v is the only vertex of W_1 , then this is obvious, as W_1 is itself a straight-ahead cycle. Thus suppose that W_1 has at least one vertex other than v . Suppose that W_1 is not a straight-ahead cycle. Then there exists a vertex u other than v that gets visited twice as we traverse W_1 . The subwalk of W_1 that starts and ends at u has fewer vertices than W_1 , and by the induction hypothesis it has a straight-ahead cycle C , which is also a straight-ahead cycle of W_1 . \square

4 Macro moves

The three Reidemeister moves are particular cases of more general operations that turn a diagram D into an *equivalent* diagram D' , that is, a diagram D' whose corresponding knot is isotopic to the knot represented by D . Following Kauffman [21], we call these operations *macro moves*. We quote the following lively passage from [21]:

[21, Section 2] “*In a macro move, we identify an arc that passes entirely under some piece of the diagram (or entirely over) and shift this part of the arc, keeping it under (or over) during the shift.*”

We also refer the reader to the interesting discussion on macro moves (although not using this terminology) in mathoverflow.net initiated by Gowers [17].

We will use two particular kinds of macro moves. Before describing them, let us clarify the notion of strand that we will use throughout this work.

We adopt the viewpoint that a *strand* of a knot K is a connected (in our context, it will always be proper) subset of K . If a diagram D is obtained as a projection of K , then we refer to the part of D that corresponds to a strand of K also as a *strand* of D . We actually only use strands in diagrams, so there should be no confusion, whenever we work with a strand, as to whether we refer to a strand as a part of a diagram or as a piece of a knot. As an illustration, we refer the reader to Figures 5 and 6. In the former, we have a strand σ whose endpoints coincide, in a crossing point, and in the latter we illustrate strands σ, σ' with distinct endpoints.

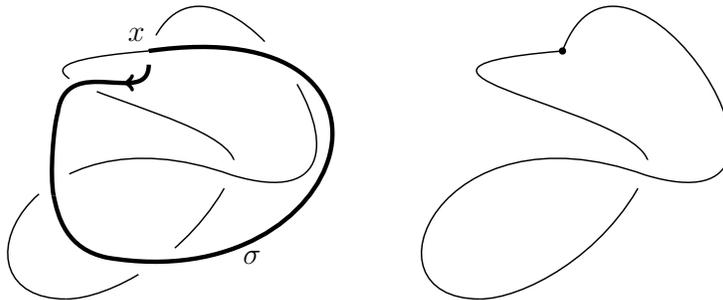


Figure 5: A macro move of Type A. In the diagram D on the left-hand side, the thick strand σ is an overstrand based at x . The diagram on the right-hand side is obtained by collapsing σ to x , that is, by removing $\sigma \setminus \{x\}$ from D . Therefore these diagrams are equivalent.

The first macro move we use is illustrated in Figure 5. Let x be a crossing point in a diagram D , and let σ be a proper strand that starts and ends at x . Further suppose that σ has no self-crossings other than x . If as we traverse σ starting from x (in any of the two possible directions), every crossing we find is an overpass (respectively, underpass) for σ , then σ is an *overstrand* (respectively, *understrand*) *based at x* .

Informally speaking, if σ is an overstrand (respectively, understrand) of D based at x , then σ lies above (respectively, below) the rest of the diagram. In either case we can then collapse σ to the point x , and the diagram D' thus obtained is equivalent to D . We say that D' is obtained from D by a *macro move of Type A*. We note that the classical Reidemeister move of Type I is a particular instance of a macro move of Type A, namely in the case in which x is the only crossing of σ .

The other macro move we use is illustrated in Figure 6. Let x, y be distinct non-crossing points in a diagram D . Suppose that σ is a strand of D with endpoints x, y , that does not cross itself, and has the property that as we traverse σ , all the crossings we find are overpasses (respectively, underpasses) for σ . Then we say that σ is an *overstrand* (respectively, *understrand*) of D . Let D' be a diagram obtained from D by removing σ , and replacing it by another strand also with endpoints x, y , that is an overstrand (respectively, understrand) in D' . Then D and D' are clearly equivalent diagrams. We say that D' is obtained from D by a *macro move of Type B*.

We note that the classical Reidemeister moves of Types II and III are particular instances of macro moves of Type B.

5 Obtaining unknot diagrams using cycle decompositions

We now formally lay out the first tool we develop to recursively construct unknot diagrams of a shadow. Let S be a shadow, and suppose that C is a straight-ahead cycle of S . Let S' be the shadow obtained from S by removing the edges of C , and suppressing any resulting degree 2 vertices. We use the notation $S' = S//C$ to denote that S' is obtained from S in this way.

The first key claim is that if $S' = S//C$, then $|\mathcal{U}(S)| \geq 2|\mathcal{U}(S')|$, and that a stronger inequality holds if C has more than one vertex:

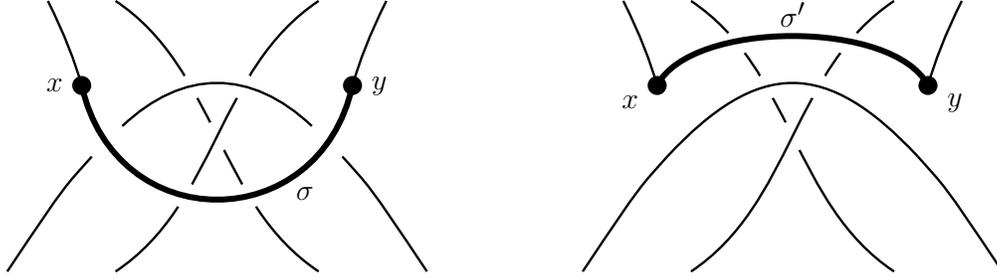


Figure 6: A macro move of Type B. The thick strand σ with endpoints x, y is an overstrand in the left-hand side diagram. We obtain the right-hand side diagram by replacing this overstrand with the overstrand σ' . These diagrams are thus equivalent.

Proposition 7. *Let S be a shadow, and let C be a straight-ahead cycle of S . Let $S' = S//C$. Then $|\mathcal{U}(S)| \geq 2|\mathcal{U}(S')|$. Moreover, if C has more than one vertex, then $|\mathcal{U}(S)| \geq 4|\mathcal{U}(S')|$.*

Proof. The main idea is that an unknot diagram of S' can be “extended” to at least two unknot diagrams of S (and to at least four, if C has more than one vertex). Suppose first that C has more than one vertex, and let c be the root of C . We refer the reader to Figure 7 for an illustration of the proof, where C is the cycle formed by the thick edges.

Starting from any unknot diagram D' of S' , we obtain four distinct unknot diagrams of S as follows. First, we let each vertex in S that is in S' maintain the prescription it has in D' . Then the prescriptions at the vertices in $S \setminus (S' \cup \{c\})$ are given so that in the resulting diagram of S , the strand corresponding to C is either an overstrand or an understrand (two possibilities). Moreover, we still have the freedom to assign any of the two possible prescriptions at c . This yields in total four distinct diagrams of S . Each of these four diagrams is unknot: one can apply a macro move of Type A to the strand corresponding to C (this is either an over- or an under-strand, so such a move can be performed), to obtain the unknot diagram D' . Thus these four diagrams are unknot.

Finally, if c is the only vertex of C , then we extend D' to two unknot diagrams of S , by giving to c each of the two possible prescriptions. \square

We wish to iteratively apply this proposition until we end up with a trivial shadow. The formal setting makes use of the following definition.

Definition. (Cycle decomposition). Let S be a shadow. Suppose that $S = S_1, S_2, \dots, S_p$ is a sequence of shadows, with the following properties:

- For $i = 1, 2, \dots, p - 1$, there is a straight-ahead cycle C_i of S_i such that $S_{i+1} = S_i//C_i$.
- S_p is the trivial shadow.

Set $C_p := S_p$. Then the sequence C_1, C_2, \dots, C_p is a *cycle decomposition* of S of size p .

In Figure 8 we give an example of a cycle decomposition. We remark that in a cycle decomposition C_1, C_2, \dots, C_p of a shadow S , the cycles C_2, \dots, C_p are not necessarily straight-ahead cycles in S . Moreover, they do not even need be cycles of S . On the other hand, for $i = 1, 2, \dots, p$ there is a cycle D_i in S such that D_i is a subdivision of C_i . The cycles D_1, D_2, \dots, D_p are pairwise edge-disjoint, and $S = \cup_{i=1}^p D_i$. The sequence D_1, D_2, \dots, D_p is the *primary sequence* of C_1, C_2, \dots, C_p .

We can now give the first main ingredient in the proof of Theorem 2.

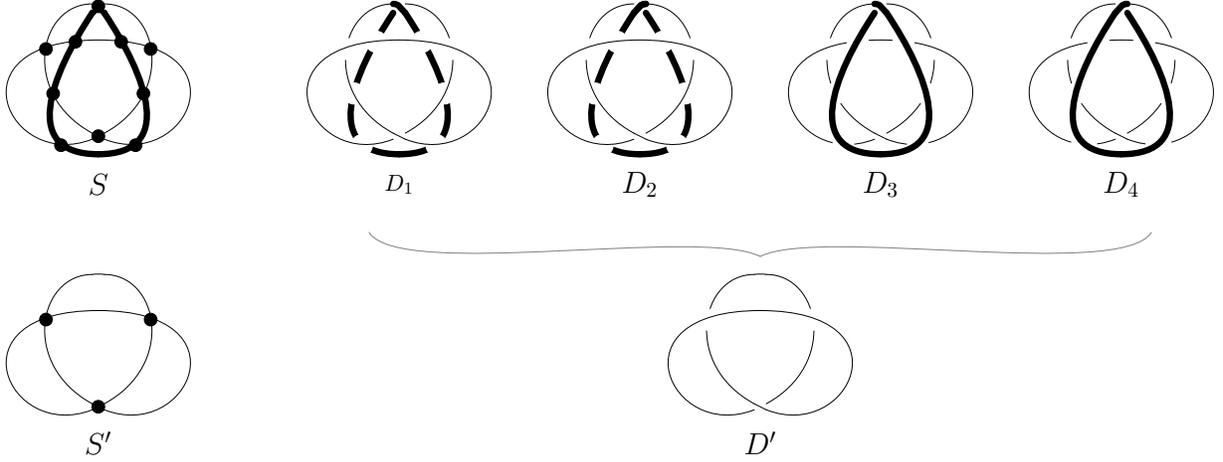


Figure 7: Illustration of the proof of Proposition 7.

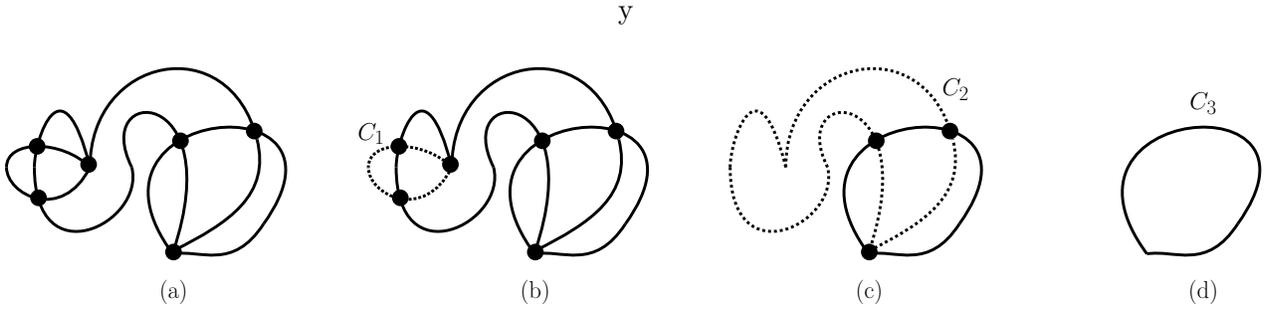


Figure 8: A cycle decomposition of the shadow $S = S_1$ in (a). In (b) we highlight with dotted edges a straight-ahead cycle C_1 of S_1 . After removing the edges of C_1 , and suppressing the three resulting degree 2 vertices, we reach the shadow $S_2 = S_1 // C_1$ in (c). The dotted cycle C_2 in (c) is a straight-ahead cycle of S_2 , and in (d) we show the shadow $S_3 = S_2 // C_2$. We finally set $C_3 = S_3$ (thus C_3 is a vertex-free cycle). The sequence C_1, C_2, C_3 is then a cycle decomposition of $S = S_1$.

Lemma 8. *Let S be a shadow with n vertices. Suppose that S has a cycle decomposition of size at least $\sqrt[3]{n}$. Then $|\mathcal{U}(S)| \geq 2^{\sqrt[3]{n}}$.*

Proof. Let S be a shadow with n vertices, and let S_1, S_2, \dots, S_p be a sequence of shadows obtained from a cycle decomposition of S , where $p \geq \sqrt[3]{n}$. A repeated application of Proposition 7 shows that $|\mathcal{U}(S)| \geq 2^{p-1} |\mathcal{U}(S_p)|$. Since $|\mathcal{U}(S_p)| = 1$ (S_p is trivial), then $|\mathcal{U}(S)| \geq 2^{p-1}$.

The bound $|\mathcal{U}(S)| \geq 2^{p-1}$ is the best we can obtain from Proposition 7 only if at each step in the cycle decomposition, we remove a cycle that has only one vertex. Otherwise we can apply at least once the second part of Proposition 7, and obtain that $|\mathcal{U}(S)| \geq 2^p$. In such a case we are then done, since $p \geq \sqrt[3]{n}$.

Thus we are done unless, at each step in the cycle decomposition, we remove a cycle that has only one vertex. It is straightforward to see that this can happen only if $p = n + 1$. In this case the bound $|\mathcal{U}(S)| \geq 2^{p-1}$ reads $|\mathcal{U}(S)| \geq 2^n$. Since $2^n > 2^{\sqrt[3]{n}}$ for every positive integer n , it follows that in this case we are also done. \square

This lemma settles Theorem 2 if S has a cycle decomposition of the given size. If all cycle decompositions of S are smaller, then we need a different approach. Before we introduce this second tool we establish some elementary results on digons in systems of two simple closed curves.

6 Digons in systems of two closed curves

Lemma 8 in the previous section settles Theorem 2 for shadows with sufficiently large cycle decompositions. For a shadow S that does not have this property, our approach will be to find two cycles B, R in S that have many vertices in common, and then to use these cycles to produce many unknot diagrams of S . This approach relies on the existence of digons (with special properties) in a system of two simple closed curves in the plane. In this section we investigate these systems of curves, and establish the required results.

Let β be a simple closed curve in the plane. A *segment* of β is a subset of β that is homeomorphic to the closed interval $[0, 1]$. Thus a segment of β is simply a non-closed curve (including its endpoints) contained in β .

Let β, ρ be two non-disjoint simple closed curves. A *digon* of (β, ρ) is a pair of segments (α, γ) , where α is in β and γ is in ρ , such that (i) α and γ have common endpoints x, y ; (ii) the only points of β that are in γ are x and y ; and (iii) the only points of ρ that are in α are x and y . In Figure 9 we illustrate this concept.

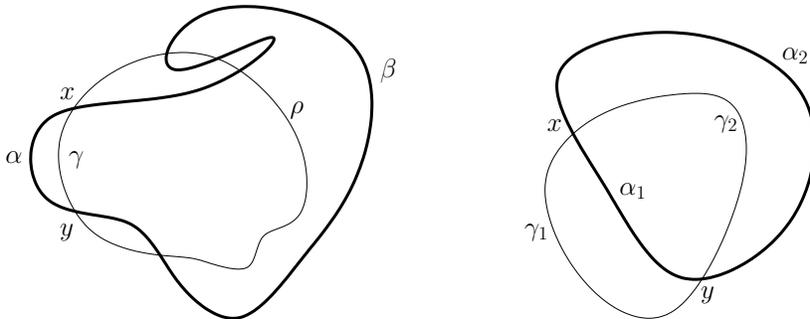


Figure 9: On the left hand side we have two simple closed curves β (thick) and ρ (thin). The pair (α, γ) is a digon of (β, ρ) . On the right hand side we have two simple closed curves $\beta = \alpha_1 \cup \alpha_2$ (thick) and $\rho = \gamma_1 \cup \gamma_2$ (thin). This last example illustrates that a digon does not necessarily bound an empty disk (that is, a disk whose interior does not contain any point in the curves). Indeed, the pair (α_2, γ_1) is a digon of $(\alpha_1 \cup \alpha_2, \gamma_1 \cup \gamma_2)$, and $\alpha_2 \cup \gamma_1$ does not bound an empty disk.

If (α, γ) is a digon of (β, ρ) , then $\alpha \cup \gamma$ is a simple closed curve. Properties (ii) and (iii) above imply that at least one of the two connected components of $\mathbb{R}^2 \setminus (\alpha \cup \gamma)$ does not have any point in $\beta \cup \rho$ (that is, at least one of these connected components is an *empty region*). We remark that it is not necessarily true that this empty region is the disk bounded by $\alpha \cup \gamma$ (see right hand side of Figure 9).

We now state and prove the results on digons that we need later. An intersection of two curves at a point x is a *crossing* if the curves intersect transversally at x . Otherwise the intersection is *tangential*.

Proposition 9. *Let β, ρ be non-disjoint simple closed curves in the plane. Suppose that every intersection point of β and ρ is a crossing (rather than tangential). Then there are at least 4*

distinct digons of (β, ρ) . Moreover, if b is a point of β that is not in ρ , and r is a point of ρ that is not in β , then there exists a digon (α, γ) of (β, ρ) such that $\alpha \cup \gamma$ contains neither b nor r .

Proof. To help comprehension, we colour β blue and ρ red. Let Δ denote the closed disk bounded by β .

We can naturally regard the part of $\beta \cup \rho$ contained in Δ as a 3-regular plane graph: each vertex is incident with two blue edges and one red edge. It is easily seen that if we take the dual of this plane graph, and discard the vertex corresponding to the unbounded face, the result is a tree. Now each leaf of this tree corresponds to a digon of (β, ρ) . Since each tree with at least two vertices has at least two leaves, then there are at least two digons of (β, ρ) whose red segments are inside Δ . A totally analogous argument shows that there are at least two digons of (β, ρ) whose red segments are outside Δ . This shows the existence of at least four distinct digons of (β, ρ) .

The existence of a digon (α, γ) of (β, ρ) such that $\alpha \cup \gamma$ contains neither b nor r is easily worked out by a simple case analysis if β and ρ have exactly two intersection points. We note that, in this case, there is exactly one digon of (β, ρ) with this property.

Suppose β and ρ have at least 4 points in common. Then if (α, γ) and (α', γ') are distinct digons, we must have $\alpha \neq \alpha'$ and $\gamma \neq \gamma'$. Thus only one digon of (β, ρ) can contain b , and also only one can contain r . As there are at least four digons of (β, ρ) , it follows that there are at least two (and in particular at least one) digons (α, γ) of (β, ρ) such that $\alpha \cup \gamma$ contains neither b nor r . \square

Proposition 10. *Let β, ρ be simple closed curves in the plane that intersect in $\ell \geq 3$ points. Suppose that there is exactly one intersection point z of β and ρ that is tangential; all the other intersection points are crossings. Then there exists a digon (α, γ) of (β, ρ) such that $\alpha \cup \gamma$ does not contain z .*

Proof. As in the proof of Proposition 9, we colour β blue and ρ red, and let Δ denote the closed disk bounded by β .

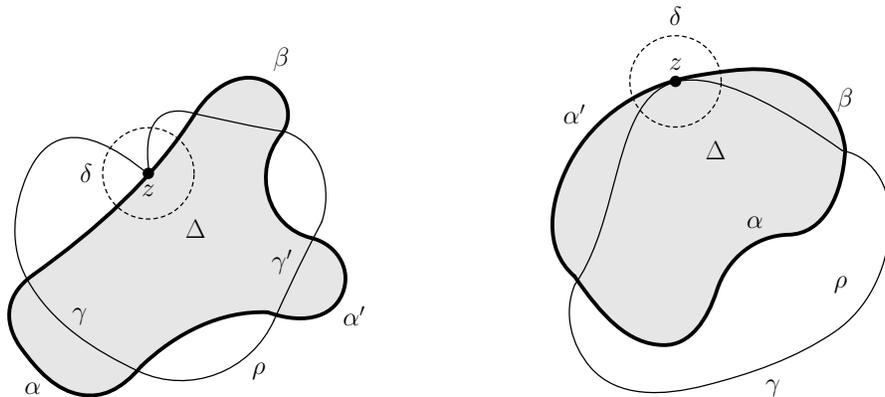


Figure 10: Illustration of the proof of Proposition 10. On the left hand side we have the case in which $\delta \cap \rho$ lies outside the disk Δ bounded by β . Both γ and γ' are segments inside Δ , and both (α, γ) and (α', γ') are digons, neither of which contains z . In the example on the right hand side, $\delta \cap \rho$ lies inside Δ , and γ is the only red segment that lies outside Δ . By letting α and α' be the blue segments into which ρ naturally breaks β , we have that (α, γ) and (α', γ) are both digons. Of these two digons, only (α, γ) does not contain z .

Suppose that if we take a small disk δ centered at z (small enough so that z is the only point in $\beta \cap \rho$ that is in δ), then $\rho \cap \delta$ lies outside Δ (see Figure 10). The other possibility ($\rho \cap \delta$ lies inside Δ) is handled in a totally analogous manner.

Now we proceed as in the proof of Proposition 9, ignoring the part of $\beta \cup \rho$ outside Δ , and find at least two distinct digons $(\alpha, \gamma), (\alpha', \gamma')$ of (β, ρ) , such that γ and γ' (which are not necessarily distinct) are inside Δ . Only one of α and α' can contain z . If α contains z , then (α', γ') is a digon of (β, ρ) that does not contain z . If α' contains z , then (α, γ) is a digon of (β, ρ) that does not contain z . \square

7 Obtaining unknot diagrams using digon-like structures

When we have a shadow S not having a sufficiently large cycle decomposition, and thus for which Lemma 8 does not guarantee enough unknot diagrams, the idea is to find two straight-ahead cycles in S with many common vertices, and then to use these two cycles to find many unknot diagrams of S . This approach is summarized in the following statement, our second main tool in the proof of Theorem 2.

Lemma 11. *Let S be a shadow, and let B and R be distinct straight-ahead cycles of S . Suppose that B and R have exactly m vertices in common. Then $|\mathcal{U}(S)| \geq 2^{m/2}$.*

Although the arguments in the proof of this lemma depend on the parity of m , the core idea is the same in both cases. We will give separately the proofs for the cases m odd and m even. Even though in principle it would have been possible to present a unified proof, we have decided against this possibility. First, it would only have resulted in a marginal saving of space. Second, dealing with the easier case (m odd) first will allow us to present the main idea in the proof, without the additional complications involved in the case when m is even.

In both cases, to help comprehension we colour the edges of B and R blue and red, respectively.

Proof of Lemma 11 for odd m . We start by noting that since m is odd, then $S = B \cup R$. Indeed, since B and R are straight-ahead cycles, it follows that if we regard them as simple closed curves, then they can have at most one non-crossing (that is, at most one tangential) intersection. An elementary argument using the Jordan curve theorem shows that the number of crossing intersections is even. Since m is odd, then these curves must have exactly one tangential intersection, which must then be the common root r of B and R . If we start a straight-ahead traversal of S starting at r , we traverse entirely one of B and R , return to r , and then traverse the other cycle, coming back again to r having traversed the whole shadow S . Thus $S = B \cup R$, as claimed.

The proof is by induction on m . The base case $m = 1$ is trivial, as in this case the lemma only requires two unknot diagrams, and it is immediately seen that a shadow with one vertex has (exactly) two unknot diagrams.

Suppose then that $m \geq 3$. It follows from Proposition 10 that there exist edges e, f with common endpoints u, v , where e is blue and f is red, such that $e \cup f$ is a digon. Moreover, it follows from that proposition that such edges e, f exist so that neither u nor v is r .

Let S' be the shadow obtained from S by “splitting” each of u and v into two degree 2 vertices, and then suppressing these four degree 2 vertices, as illustrated in Figure 11. It is easy to see that the resulting graph S' is indeed a shadow, with straight-ahead cycles B', R' naturally induced from B and R , whose common root is still r , and such that $S' = B' \cup R'$.

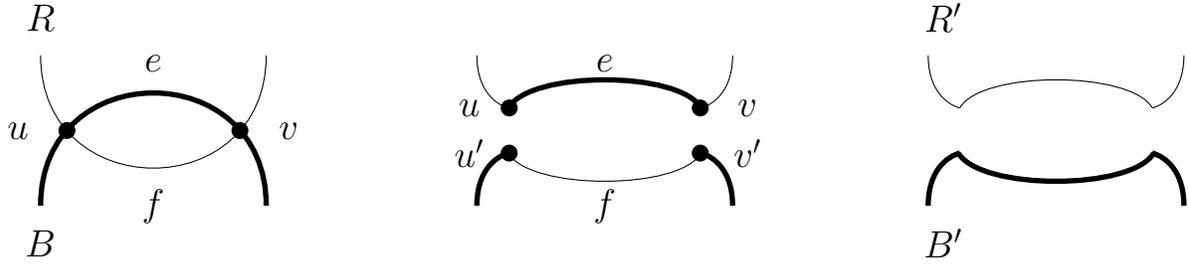


Figure 11: Obtaining S' from S by splitting two vertices. The relevant part of the blue cycle B (respectively, red cycle R) is drawn with thick (respectively, thin) edges. The blue edge e and the red edge f have common endpoints u, v . We split each of u and v into two degree 2 vertices, and then suppress these four vertices. The resulting shadow S' also has the property that it is the union of two straight-ahead cycles. We can naturally label these cycles B' and R' , letting B' (respectively, R') be the cycle that contains the remains of the red edge f (respectively, the blue edge e).

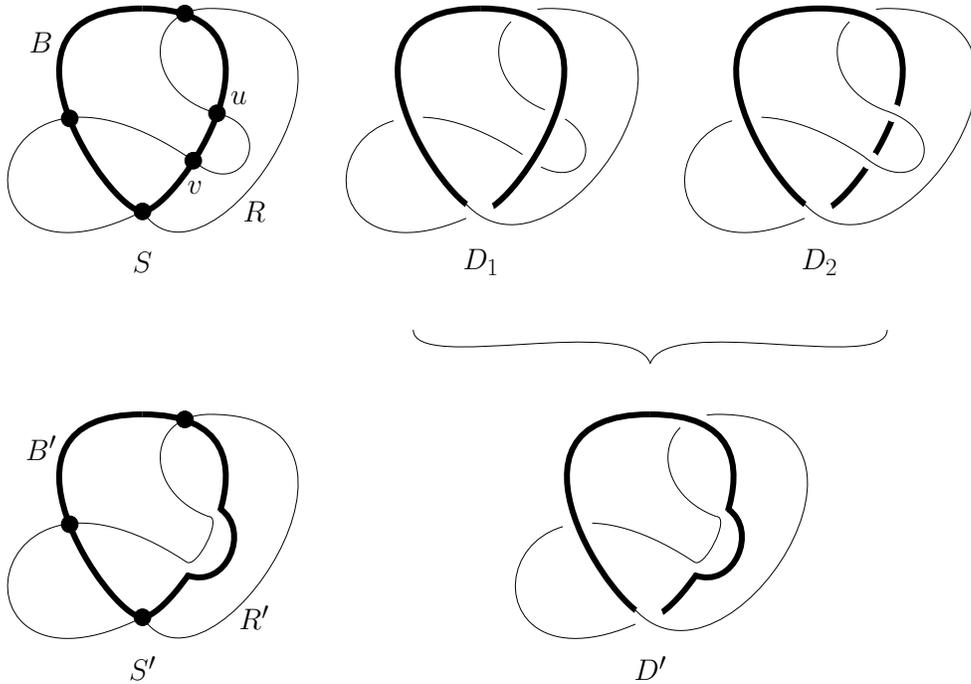


Figure 12: The shadow S' is obtained from S by splitting u and v as in Figure 11. After performing this operation, the straight-ahead cycles B and R of S naturally induce straight-ahead cycles B' and R' of S' . Any unknot diagram of S' can be extended to two unknot diagrams of S . For instance, the illustrated diagram D' of S' can be extended to the two unknot diagrams D_1, D_2 shown. In the digon with endcrossings u, v in D_1 , the strand corresponding to B is an overstrand, and in D_2 it is an understrand. In either case, a Reidemeister move of Type II can be used to transform the diagram into D' . Since D' is unknot, then D_1 and D_2 are also unknot.

The key claim is that every unknot diagram of S' can be extended to two unknot diagrams of S . See Figure 12 for an illustration. For let D' be an unknot diagram of S' . We construct from D' two diagrams D_1, D_2 as follows. For each vertex in S that is in S' , we maintain the prescription it has in D' . We obtain D_1 (respectively, D_2) by giving prescriptions to u and v such that the corresponding crossings are overpasses for the blue (respectively, red) strand.

Both D_1 and D_2 are unknot diagrams of S . This follows simply since, by applying a Reidemeister move of Type II to the strands corresponding to e and f , we obtain the unknot diagram D' . Thus D_1 and D_2 are unknot.

This argument shows that S has at least twice as many unknot diagrams as S' , and so the lemma follows by induction. \square

Proof of Lemma 11 for even m . A similar reasoning as the one at the beginning of the proof of Lemma 11 shows that in this case necessarily $S \neq B \cup R$, and that the respective roots b and r of B and R must be distinct. Since $S \neq B \cup R$, then S has edges that are neither blue nor red; we colour these edges gray.

Let \bar{S} denote the graph that results by removing from S all edges not in $B \cup R$, and suppressing all resulting degree 2 vertices. Thus \bar{S} is the shadow of a link with two components: \bar{S} has exactly two straight-ahead cycles, namely the cycles \bar{B} and \bar{R} naturally inherited from B and R .

Note that the roots b and r are not anymore vertices in \bar{S} , but we can still regard them as middle points \bar{b} and \bar{r} of edges of \bar{B} and \bar{R} , respectively. We say that an unlink diagram of \bar{S} is (\bar{b}, \bar{r}) -avoiding if it can be taken to the trivial unlink diagram by a series of Reidemeister moves of Type II, none of which involves moving or passing above/below neither \bar{b} nor \bar{r} .

It is easy to see that arguments totally analogous to those in the proof above for the case m odd (but in the present case using Proposition 9 instead of Proposition 10) yields that \bar{S} has at least $2^{m/2}$ distinct (\bar{b}, \bar{r}) -avoiding unlink diagrams.

The key point that we will now show is that each of these $2^{m/2}$ (\bar{b}, \bar{r}) -avoiding unlink diagram of \bar{S} can be extended to an unknot diagram of S . Given an (\bar{b}, \bar{r}) -avoiding unlink diagram \bar{D} of \bar{S} , we extend it to a diagram D of S as follows:

- (i) For each crossing of D that is in \bar{D} , we let it maintain in D the prescription it has in \bar{D} .
- (ii) If a crossing in D involves a blue (respectively, red) strand and a gray strand, then we let it be an overpass for the blue (respectively, red) strand. Loosely speaking, the blue and red strands are “above” the gray strands.
- (iii) For the vertices that are incident with four gray edges (if any) we proceed as follows. Let S' be the shadow that results by removing the edges of $B \cup R$, and suppressing all resulting degree 2 vertices. We give the prescription on the vertices of S' by asking that the resulting assignment on S' is an unknot diagram (take for instance any descending diagram of S').
- (iv) We give arbitrary prescriptions to the crossings corresponding to b and r .

These rules give prescriptions to all the vertices in S . We claim that the resulting diagram D of S is unknot.

To prove that D is unknot, we start by arguing that the blue-red crossings can be eliminated by using macro moves of Type B. For consider a sequence of (\bar{b}, \bar{r}) -avoiding Reidemeister Type II moves that unlink \bar{D} . Since the blue and red strands are “above” the gray strands (by condition

(ii)), each of this moves can be translated to a macro move of Type B on D ; this is the step in which it is essential that these moves on \overline{D} are $(\overline{b}, \overline{r})$ -avoiding.

Thus we can eliminate from D all blue-red crossings by a sequence of macro moves of Type B. After this, the remaining blue and red strands are still overstrands (again by condition (ii)), and so they can be collapsed to \overline{b} and \overline{r} , respectively, by using macro moves of Type A. Since by (iii) the resulting diagram is unknot, it follows that D is an unknot diagram, as claimed.

Thus each $(\overline{b}, \overline{r})$ -avoiding unlink diagram of \overline{S} can be extended to an unknot diagram of S . Since these extensions are all distinct to each other, and the number of $(\overline{b}, \overline{r})$ -avoiding unlink diagrams is at least $2^{m/2}$, it follows that $|\mathcal{U}(S)| \geq 2^{m/2}$. \square

8 Proof of Theorem 2

Lemmas 8 and 11 take care of Theorem 2 when S is a shadow that either (i) has a cycle decomposition of size at least $\sqrt[3]{n}$; or (ii) has two straight-ahead cycles with at least $2\sqrt[3]{n}$ common vertices.

Unfortunately, there exist shadows that satisfy neither (i) nor (ii). However, we can show that if S does not satisfy (i), then there is a substructure of S that does satisfy (ii). To make this precise, let us introduce the concept of a subshadow. A shadow T is a *subshadow* of S if there is a sequence $S = S_1, S_2, \dots, S_r = T$ of shadows, such that for $i = 1, 2, \dots, r - 1$, there is a straight-ahead cycle C_i of S_i such that $S_{i+1} = S_i // C_i$. Note that a shadow is a subshadow of itself.

An easy but crucial fact that follows from Proposition 7 is that if T is a subshadow of S , then $|\mathcal{U}(S)| \geq |\mathcal{U}(T)|$. Thus in order to apply Lemma 11 to obtain Theorem 2, we do not need the existence of two straight-ahead cycles with these properties in S itself; it suffices to guarantee the existence of such straight-ahead cycles in some subshadow T of S . This last ingredient is provided by the following statement.

Lemma 12. *Let S be a shadow with n vertices. Suppose that every cycle decomposition of S has size at most $\sqrt[3]{n}$. Then there is a subshadow T of S that has straight-ahead cycles B and R , such that B and R have at least $2\sqrt[3]{n}$ vertices in common.*

Before moving on to the proof of this lemma, for completeness let us formally show that with this last ingredient we can finally prove Theorem 2.

Proof of Theorem 2. If S has a cycle decomposition of size at least $\sqrt[3]{n}$, then we are done by Lemma 8. If every cycle decomposition of S has size at most $\sqrt[3]{n}$, then by Lemma 12 it follows that S has a subshadow T with straight-ahead cycles B and R that have at least $2\sqrt[3]{n}$ common vertices. Applying Lemma 11 to T , we obtain $|\mathcal{U}(T)| \geq 2\sqrt[3]{n}$. Since $|\mathcal{U}(S)| \geq |\mathcal{U}(T)|$, the theorem follows. \square

The proof of Lemma 12 relies on the following two results on cycle decompositions.

Proposition 13. *Let S be a nontrivial shadow, and let C be a straight-ahead cycle of S . Then there is at most one straight-ahead cycle of $S // C$ that is not a subdivision of a straight-ahead cycle of S .*

Proof. Let $S' := S // C$, and let \mathcal{F}' be the set of all straight-ahead cycles of S' . Each $F' \in \mathcal{F}'$ is obtained from a cycle F in S , by suppressing from F the vertices that are in C (if any). Let \mathcal{F}

denote the family of all such cycles F . It is readily seen that each $F \in \mathcal{F}$ is a straight-ahead cycle in S , unless F contains the root of C . Since at most one element in \mathcal{F} can contain the root of C , the proposition follows. \square

Proposition 14. *Let S be a nontrivial shadow. Let $\mathcal{C} = C_1, C_2, \dots, C_p$ be a cycle decomposition of S , and let D_1, D_2, \dots, D_p be the primary sequence of \mathcal{C} . Then at least two elements in $\{D_1, D_2, \dots, D_p\}$ are straight-ahead cycles of S .*

Proof. We proceed by induction on p . The base case $p = 2$ is straightforward. Thus we take a cycle decomposition C_1, C_2, \dots, C_p of a shadow S , with $p \geq 3$, and assume that the statement holds for all cycle decompositions (of any shadow) of size less than p .

Evidently, C_1 is a straight-ahead cycle of S . Let $S_2 := S // C_1$. Consider the cycle decomposition $\mathcal{C}_2 = C_2, C_3, \dots, C_p$ of S_2 , and let E_2, E_3, \dots, E_p be the associated primary sequence. By the induction hypothesis, there exist distinct $j, k \in \{2, \dots, p\}$ such that E_j, E_k are straight-ahead cycles of S_2 .

By Proposition 13, at most one of E_j and E_k is not a subdivision of a straight-ahead cycle of S . Without loss of generality, E_j is a subdivision of a straight-ahead cycle of S . Now D_j is a subdivision of E_j , and so it follows that D_j is a straight-ahead cycle in S . This completes the proof, since D_1 and D_j are straight-ahead cycles in S . \square

Proof of Lemma 12. Let S be a shadow with n vertices, and let C_1, C_2, \dots, C_p be a cycle decomposition of S , with $p \leq \sqrt[3]{n}$. Let D_1, D_2, \dots, D_p be the primary sequence of C_1, C_2, \dots, C_p .

We start by noting that there exist $r, s \in \{1, 2, \dots, p\}$ such that D_r and D_s have at least $2\sqrt[3]{n}$ common vertices. To see this, note that each vertex in S is in the intersection of exactly two cycles in D_1, D_2, \dots, D_p . If each pair of cycles in this sequence had fewer than $2\sqrt[3]{n}$ common vertices, this would imply that S has fewer than $\binom{p}{2} \cdot 2\sqrt[3]{n} \leq \binom{\sqrt[3]{n}}{2} \cdot 2\sqrt[3]{n} = n - n^{2/3}$ vertices.

Now let T be a subshadow of S such that (i) T has a cycle decomposition E_1, E_2, \dots, E_q , with primary sequence F_1, F_2, \dots, F_q such that both D_r and D_s are subdivisions of cycles in $\{F_1, F_2, \dots, F_q\}$; and (ii) no subshadow of S with fewer vertices than T satisfies (i). Note that T is well defined, as S itself satisfies (i), as witnessed by its cycle decomposition C_1, C_2, \dots, C_p . We claim that T satisfies the conditions in the lemma.

To show this, let j, k be the integers in $\{1, 2, \dots, q\}$ such that D_r is a subdivision of F_j , and D_s is a subdivision of F_k .

We proceed by contradiction. Suppose that one of F_j and F_k is not a straight-ahead cycle of T . Then, by Proposition 14, there is an $\ell \in \{1, 2, \dots, q\}$, $\ell \neq j, k$, such that F_ℓ is a straight-ahead cycle of T . Let $T' := T // F_\ell$. Then E_1, E_2, \dots, E_q naturally induces a cycle decomposition $E'_1, E'_2, \dots, E'_{\ell-1}, E'_{\ell+1}, \dots, E'_q$ of T' . Indeed, for each $i \in \{1, 2, \dots, \ell-1, \ell+1, \dots, q\}$ it suffices to let E'_i be the cycle that results by suppressing from E_i all the vertices that are in E_ℓ , if any.

Let $F'_1, F'_2, \dots, F'_{\ell-1}, F'_{\ell+1}, \dots, F'_q$ be the primary sequence of $E'_1, E'_2, \dots, E'_{\ell-1}, E'_{\ell+1}, \dots, E'_q$. Then, for each $i \in \{1, 2, \dots, \ell-1, \ell+1, \dots, q\}$ we have that F_i is a subdivision of F'_i . In particular, F_j is a subdivision of F'_j , and F_k is a subdivision of F'_k . It follows that back in S , D_r is a subdivision of F'_j , and D_s is a subdivision of F'_k . Since this violates the minimality of T , we conclude that F_j and F_k are straight-ahead cycles of T . The lemma now follows by setting $B = F_j$ and $R = F_k$, since these two cycles have the same number of common vertices as D_r and D_s , which is at least $2\sqrt[3]{n}$. \square

9 Proofs of Theorem 3 and Observation 4

Let us start by recalling the concept of a cut-vertex in a graph. Following Tutte [29], a *1-separation* of a connected graph G is an ordered pair (H, K) of subgraphs of G , each with at least one edge, such that $H \cup K = G$ and $H \cap K$ is a graph that consists of a single vertex. The vertex of $H \cap K$ is the *cut-vertex* of the 1-separation.

We remark that sometimes a cut-vertex of a connected graph is thought of as a vertex whose deletion produces a disconnected graph. We note that Tutte's notion of a cut-vertex is more general. Consider for instance vertex v in the shadow at the top of Figure 2. If we delete this vertex, together with its incident edges, the resulting graph is still connected, and yet v is a cut-vertex according to Tutte's definition. Under Tutte's definition, if a graph G can be obtained from two disjoint graphs, taking one vertex of each of these graphs and identifying these two vertices into a vertex v , then v is a cut-vertex of G .

9.1 Proof of Theorem 3

We prove the two statements in Theorem 3 separately.

(I) *Every diagram of S is unknot if and only if every vertex of S is a cut-vertex.*

For brevity, let us say that a shadow is *simple* if each of its vertices is a cut-vertex.

For the easier direction of (I), we note that if D is a diagram of a simple shadow, then every crossing in D is nugatory. We recall that a *nugatory* crossing in a diagram corresponds precisely to a cut-vertex in its shadow. This crossing can be trivialized by a switching operation of the diagram, thus obtaining an equivalent diagram with fewer crossings. If every crossing in a diagram is nugatory, then it is equivalent to the trivial diagram. Thus, if S is simple, then all the diagrams of S are unknot.

We prove the other direction of (I) by induction on the number of vertices of S . In the base case S consists of exactly one vertex, and the statement is trivial. For the inductive step, we assume that the statement holds for all simple shadows with fewer than n vertices, for some $n \geq 2$, and let S be a simple shadow with n vertices.

First we note that if S has no cut-vertex, then not all the diagrams of S are unknot. To see this, let D be an alternating diagram of S , that is, a diagram in which as one traverses the diagram, one encounters an overcrossing, then an undercrossing, then an overcrossing, and so on. Since S has no cut vertices, then D has no nugatory crossings, that is, D is a *reduced* diagram. The validity [19, 23, 28] of the First Tait Conjecture (if D is a reduced alternating diagram, then no diagram equivalent to D has fewer crossings than D) then implies that the knot represented by D cannot be the unknot.

We can then assume that S has a cut-vertex v . Then there are subgraphs H, K of S , both with at least one edge, such that $H \cup K = S$ and $H \cap K$ consists solely of v . Note that v has degree 2 in both H and K . Let H' (respectively, K') be the graph obtained by suppressing v in H (respectively, K). Then H' and K' are shadows, each with fewer than n vertices.

We claim that both H' and K' are simple. For suppose that H' is not simple, and let D_H be a knotted diagram of H' . Let D_K be an unknot diagram of K' . We can naturally combine D_H and D_K into a diagram D of S , by giving any prescription to v . Since D_K is unknot, it follows that D and D_H are equivalent, and so D is a knotted diagram of S . Since this contradicts the assumption that S is simple, we conclude that H' must be simple. The same argument shows that K' is simple.

The inductive hypothesis then implies that every vertex of H' and K' is a cut-vertex, and from this it immediately follows that every vertex of S is a cut-vertex. \square

(II) *If not every vertex of S is a cut-vertex, then S has a trefoil diagram.*

It is easy to see that if every straight-ahead cycle of a shadow has only one vertex (that is, it consists of a vertex and a loop-edge), then every vertex is a cut-vertex of the shadow. From (I) it then follows that if not all the vertices of a shadow are cut vertices, then the shadow has a straight-ahead cycle with more than one vertex.

The proof is by induction on the number of vertices of S . A quick analysis shows that in a shadow with 1 or 2 vertices, all the vertices are cut vertices. It is also readily seen that up to isomorphism there is only one shadow with 3 vertices, such that not all of its vertices are cut vertices. This is the shadow of the familiar diagram of the trefoil knot. This shadow has evidently a trefoil diagram, and so this settles the base case $n = 3$.

For the inductive step, we take a shadow S on $n \geq 4$ vertices, where not all the vertices of S are cut vertices, and assume that (II) holds for shadows with fewer than n vertices.

Let S be a shadow in which not every vertex of S is a cut-vertex. It follows from the observation above that S has a straight-ahead cycle with more than one vertex. Let C be such a straight-ahead cycle, and let r be its root. Let W be the other straight-ahead closed walk that starts and ends at r . Note that $S = C \cup W$.

The walk W can then be naturally written as the concatenation of three walks W_1, W_2 , and W_3 , as follows. Since C has at least one vertex other than r , and W is a closed walk and $S = C \cup W$, then W and C must have at least two vertices in common, other than r .

Let u be the first vertex of C (other than r) that we find as we traverse W . After reaching u , we continue traversing W ; let v be the first vertex in C that we find in this continuation of the traversal of W . Then let W_1 be the part of W from r to u , let W_2 be the part of W from u to v , and let W_3 be the part from v to r . Note that W_1 is contained in a component Δ_1 of $\mathbb{R}^2 \setminus C$, and W_2 is contained in the other component Δ_2 . The end edges of W_3 are contained in Δ_1 , but some parts of W_3 may be in Δ_1 .

The situation is illustrated in Figure 13(a). The dashed parts of W_1 and W_2 hint to the fact that W_1 and W_2 may not be paths. The dashed parts of W_3 suggest that the beginning and the end of W_3 are in Δ_1 , but other parts of W_3 may be in Δ_2 .

We now show that if one of W_1, W_2, W_3 is not a path, then we are done by the induction hypothesis. For suppose that W_1 is not a path. Since W_1 is a straight ahead-walk, it follows that there must exist a straight-ahead cycle F contained in W_1 . We note that none of r, u , and v can be in F . Now let $T := S // F$. The shadow T has then a straight-ahead cycle (namely the one naturally induced by C) with more than one vertex, and so by the induction hypothesis T has a trefoil diagram. But this trefoil diagram can be extended to an equivalent diagram of S , by having the strand corresponding to F to be an overstrand. Therefore we conclude that S has a trefoil diagram.

In exactly the same way it is shown that if one of W_2 and W_3 is not a path, then we are done by the induction hypothesis. We can then assume that W_1, W_2 , and W_3 are paths. A possible scenario is illustrated in Figure 13(b). We finish the proof by showing that S has a trefoil diagram.

Consider a diagram D of S in which the strand corresponding to W_3 is an overstrand. An example is illustrated in Figure 14(b), which is a diagram of the shadow in Figure 13.

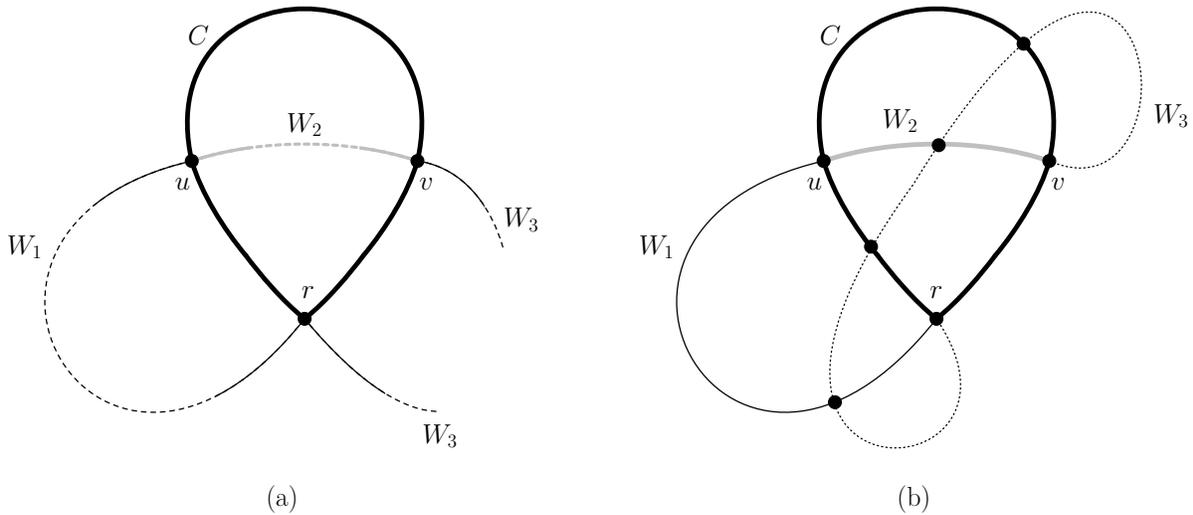


Figure 13: If S is a shadow with a straight-ahead cycle C that has more than one vertex, then S must have roughly the structure shown in (a). That is, the complement of C in S can be naturally decomposed into three parts W_1, W_2 and W_3 , where W_1 and W_2 are contained in distinct components of $\mathbb{R}^2 \setminus C$, and W_3 need not be contained in any of these components, but its end edges are contained in the same component as W_1 . In part (b) of this figure we depict a whole shadow that satisfies these properties.

With a macro move of Type B we can replace this strand with a strand that is contained in Δ_2 , and has no internal crossings. Let D' be the diagram thus obtained. (See Figure 14(b)). The shadow of D' is then isomorphic to the usual shadow of the trefoil knot, and so it follows that by giving suitable prescriptions at r, u , and v (as in Figure 14(b)), we obtain that D' is a diagram of the trefoil knot. By giving these prescriptions at r, u , and v also in D , it follows that D and D' are equivalent, and so D is a trefoil diagram of S . \square

9.2 Proof of Observation 4

For each integer $n \geq 3$, let C_n be the shadow obtained by taking a cycle with n vertices and adding a parallel edge to each edge. In Figure 15 we illustrate C_7 .

To prove Observation 4 it suffices to show that for every odd integer n , if D is a diagram of C_n , then the knot corresponding to D is either trivial, or it has odd crossing number.

We prove this by induction on n . For the base case $n = 3$ it is readily seen that every diagram of C_3 is either a trefoil diagram or an unknot diagram. Suppose then that $n \geq 5$, and that the statement holds for all odd integers smaller than n .

Let D be a diagram of C_n . Since C_n is a 2-connected graph, it follows that D is a reduced diagram. If D is alternating then, since D is reduced, the crossing number of the knot corresponding to D is n [19, 23, 28]. Suppose finally that D is not alternating. Then there are two neighbour crossings on which we can perform a Reidemeister move of Type II on D , obtaining a diagram of

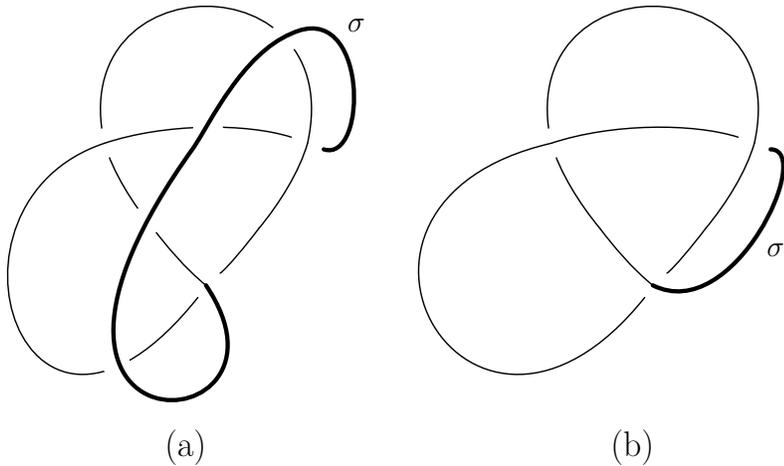


Figure 14: On the left hand side we have a diagram of a shadow in which W_1, W_2 , and W_3 are paths (this is actually a diagram of the shadow in Figure 13(b)), and the strand σ corresponding to W_3 is an overstrand. We can use a macro move of Type B to shift σ to a strand σ' with no internal crossings, to obtain the equivalent diagram shown in (b). This last diagram is clearly a trefoil diagram, and so the diagram in (a) is a trefoil diagram of the shadow in Figure 13(b).

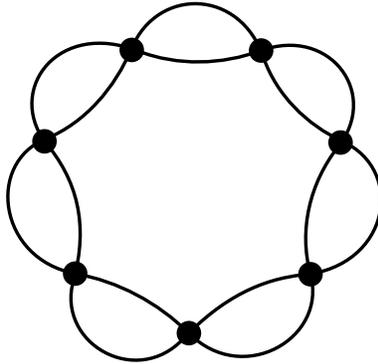


Figure 15: The shadow C_7 . No diagram associated to C_7 is equivalent to a diagram of a knot with even crossing number (and in particular to the figure-eight knot).

C_{n-2} , and the result follows by the induction hypothesis. □

10 Concluding remarks and open questions

Before starting this project, we considered the following question. Given a shadow S on n vertices, how many different knot types are there among the 2^n distinct diagrams of S ? This question was posed by the third author in mathoverflow.net [26]. Among the replies received, Andy Putnam

mentioned shadows that have only unknot diagrams. After obtaining the easy characterization of which shadows have this property (the first part of Theorem 3), we turned our attention to the main topic in this paper.

We find the question in the previous paragraph quite interesting. We leave aside those shadows that only have unknot diagrams, as they admit an easy characterization. The graphs used in the proof of Observation 4 show that for each odd integer n , there is a 2-connected shadow C_n on n vertices, such that among the 2^n distinct diagrams of C_n there are only $\lceil n/2 \rceil$ distinct knot types. Now, even though very large, these shadows have very small depth [14].

We adapt the notion of depth of a simple curve, to the parallel context of the depth of a shadow. The *depth* of a face f of a shadow S is the length of a minimum distance path, in the dual graph of S , from the dual vertex of f to the dual vertex of the unbounded face. The *depth* of S is then the maximum depth over all faces of S .

Is it true that one can guarantee many more knot types for 2-connected shadows with large depth? We note that the 2-connectedness requirement is important, as the shadow on top of Figure 2 can be generalized to obtain a shadow with arbitrarily large depth, and that only has unknot diagrams. Perhaps some of the tools developed in [14] can shed light on this question.

It seems natural to ask if one could hope to prove a much better bound than the one given by Theorem 2. To explore this question, we let $\mathcal{U}_{\min}(n)$ denote the best possible bound. That is, for each positive integer n ,

$$\mathcal{U}_{\min}(n) := \min_{|V(S)|=n} |\mathcal{U}(S)|,$$

where the minimum is taken over all shadows with n vertices.

Under this terminology, Theorem 2 gives the superpolynomial bound $\mathcal{U}_{\min}(n) \geq 2^{\sqrt[3]{n}}$. Leaving aside relatively marginal possible improvements, such as showing that $\mathcal{U}_{\min}(n) \geq 2^{\sqrt{n}}$, the important question seems to be the following.

Question 15. *Is $\mathcal{U}_{\min}(n)$ an exponential, or a subexponential function?*

It follows from the work in [6, 7] that there is a constant d , $1 < d < 2$, such that $\mathcal{U}_{\min}(n) < d^n$. To find explicit upper bounds for d , we can proceed as follows. Consider any fixed shadow T with m vertices, and let $t := |\mathcal{U}(T)|$. Now let T^k denote the connected sum T^k of k copies of T . (One can easily extend to shadows the definition of the connected sum of diagrams). Since the connected sum of two diagrams is unknot if and only if each of the diagrams is unknot, it follows that T^k has exactly t^k unknot diagrams. Thus T^k is a shadow with $n := km$ vertices and $|\mathcal{U}(T^k)| = t^k = (t^{k/n})^n = (t^{1/m})^n$. From this it follows that $\mathcal{U}_{\min}(n) \leq t^{1/m}$.

To apply this approach we have explored a variety of shadows for which we can estimate (upper bound) their number of unknot diagrams. Using SnapPy [8] we found that there is a shadow S on 16 vertices such that $|\mathcal{U}(S)| < 6416$. From the previous discussion it follows that, for every positive integer n divisible by 16, $\mathcal{U}_{\min}(n) \leq (6416^{1/16})^n \approx (1.729)^n$.

Theorem 3 gives a structural characterization of which shadows have a trefoil diagram, and Observation 4 shows that there are arbitrarily large 2-connected shadows that do not have a figure-eight diagram. It seems natural to inquire what can be said about the figure-eight knot:

Question 16. *Is there a simple characterization of which shadows have a figure-eight diagram? Is it true that every sufficiently large 3-connected shadow contains a diagram of the figure-eight knot?*

Acknowledgements

This work started while the second author visited the other authors, supported by the Laboratorio Internacional Asociado Solomon Lefschetz (LAISLA), which has now become the Unité Mixte Internationale CNRS-CONACYT-UNAM “Laboratoire Solomon Lefschetz”. The first and third author were supported by Conacyt grant 222667. We thank Ronald Ortner for helpful comments, and Mario Eudave for his guidance at the beginning of this project.

References

- [1] Colin C. Adams, *The knot book. An elementary introduction to the mathematical theory of knots*, American Mathematical Society, Providence, RI, 2004.
- [2] Colin Adams, Thomas Crawford, Benjamin DeMeo, Michael Landry, Alex Tong Lin, MurphyKate Montee, Seojung Park, Saraswathi Venkatesh, and Farrah Yhee, *Knot projections with a single multi-crossing*, *J. Knot Theory Ramifications* **24** (2015), no. 3, 1550011, 30 pp.
- [3] J. Arsuaga, T. Blackstone, Y. Diao, E. Karadayi, and M. Saito, *Linking of uniform random polygons in confined spaces*, *J. Phys. A* **40** (2007), no. 9, 1925–1936.
- [4] Gregory R. Buck, *Random knots and energy: elementary considerations*, *J. Knot Theory Ramifications* **3** (1994), no. 3, 355–363. *Random knotting and linking* (Vancouver, BC, 1993).
- [5] Jason Cantarella, Harrison Chapman, and Matt Mastin, *Knot probabilities in random diagrams*, *J. Phys. A* **49** (2016), no. 40, 405001, 28pp.
- [6] Harrison Chapman, *Asymptotic laws for knot diagrams*, *Disc. Math. and Theoret. Comp. Sci. proc.* **BC** (2016), 323–334. *FPSAC 2016* (Vancouver, Canada).
- [7] ———, *Asymptotic laws for random knot diagrams*, *J. Phys. A* **50** (2017), no. 22, 225001, 32pp.
- [8] Marc Culler, Nathan M. Dunfield, Matthias Goerner, and Jeffrey R. Weeks, *SnapPy, a computer program for studying the geometry and topology of 3-manifolds*. Available at <http://snappy.computop.org> (17/02/2017).
- [9] M. Delbruck, *Knotting problems in biology*, Richard Bellman, editor. *Mathematical Problems in the Biological Sciences. Proceedings of Symposia in Applied Mathematics.* 14, American Mathematical Society, 1962, pp. 55–63.
- [10] Yuanan Diao, Claus Ernst, and Uta Ziegler, *Generating large random knot projections*, *Physical and numerical models in knot theory*, Ser. Knots Everything, vol. 36, World Sci. Publ., Singapore, 2005, pp. 473–494.
- [11] Y. Diao, J. C. Nardo, and Y. Sun, *Global knotting in equilateral random polygons*, *J. Knot Theory Ramifications* **10** (2001), no. 4, 597–607.
- [12] Yuanan Diao, Nicholas Pippenger, and De Witt Sumners, *On random knots*, *J. Knot Theory Ramifications* **3** (1994), no. 3, 419–429. *Random knotting and linking* (Vancouver, BC, 1993).
- [13] N. Dunfield, A. Hirani, M. Obeidin, A. Ehrenberg, S. Bhattacharyya, and D. Lei, *Random Knots: A preliminary report, 2014*. (2014). http://www.math.uiuc.edu/~nmd/preprints/slides/random_knots.pdf.
- [14] Hsien-Chih Chang and Jeff Erickson, *Untangling planar curves*, 32nd International Symposium on Computational Geometry, LIPIcs. Leibniz Int. Proc. Inform., vol. 51, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2016, pp. Art. 29, 16.
- [15] Chaim Even-Zohar, Joel Hass, Nati Linial, and Tahl Nowik, *Invariants of random knots and links*, *Discrete Comput. Geom.* **56** (2016), no. 2, 274–314.
- [16] H.L. Frisch and E. Wasserman, *Chemical topology 1*, *J. Am. Chem. Soc.* **83(18)** (1961), 3789–3795.
- [17] Timothy Gowers (<http://mathoverflow.net/users/1459/gowers>), *Are there any very hard unknots?*, *MathOverflow*. URL: <http://mathoverflow.net/q/53471> (version: 2014-07-29).
- [18] Douglas Jungreis, *Gaussian random polygons are globally knotted*, *J. Knot Theory Ramifications* **3** (1994), no. 4, 455–464.
- [19] Louis H. Kauffman, *State models and the Jones polynomial*, *Topology* **26** (1987), no. 3, 395–407.

- [20] ———, *Knot diagrammatics*, Handbook of knot theory, Elsevier B. V., Amsterdam, 2005, pp. 233–318.
- [21] ———, *The unknotting problem*, Open problems in mathematics, Springer, [Cham], 2016, pp. 303–345.
- [22] Maciej Mroczkowski, *Diagrammatic unknotting of knots and links in the projective space*, J. Knot Theory Ramifications **12** (2003), no. 5, 637–651.
- [23] Kunio Murasugi, *Jones polynomials and classical conjectures in knot theory*, Topology **26** (1987), no. 2, 187–194.
- [24] Makoto Ozawa, *Ascending number of knots and links*, J. Knot Theory Ramifications **19** (2010), no. 1, 15–25.
- [25] Nicholas Pippenger, *Knots in random walks*, Discrete Appl. Math. **25** (1989), no. 3, 273–278.
- [26] Gelasio Salazar (<https://mathoverflow.net/users/64175/gelasio-salazar>), *How many different knot types can have the same shadow (projection)*, MathOverflow. URL: <https://mathoverflow.net/q/250603> (version: 2017-04-13).
- [27] D. W. Sumners and S. G. Whittington, *Knots in self-avoiding walks*, J. Phys. A **21** (1988), no. 7, 1689–1694.
- [28] Morwen B. Thistlethwaite, *A spanning tree expansion of the Jones polynomial*, Topology **26** (1987), no. 3, 297–309.
- [29] W. T. Tutte, *Connectivity in graphs*, Mathematical Expositions, No. 15, University of Toronto Press, Toronto, Ont.; Oxford University Press, London, 1966.