msp

Algebraic & Geometric Topology 20 (2020) 643–698

The Reidemeister graph is a complete knot invariant

AGNESE BARBENSI
DANIELE CELORIA

We describe two locally finite graphs naturally associated to each knot type K, called *Reidemeister graphs*. We determine several local and global properties of these graphs and prove that in one case the graph-isomorphism type is a complete knot invariant up to mirroring. Lastly, we introduce another object, relating the Reidemeister and Gordian graphs, and determine some of its properties.

57M25

1 Introduction

The *Gordian graph* is a well-known graph in knot theory; its vertices are given by knot types, and two knots have an edge between them whenever they are related by a single crossing change. This graph can be thought of as describing knot theory at "large scales".

The Gordian graph is, however, very ill behaved: Each vertex of this graph has infinite valence, and vertices at distance 2 are connected by infinitely many distinct minimal paths; see Baader [3]. For every $n \geq 1$ there are embeddings of the graphs $\mathbb{Z}^n$ into it; see Marché [14].

This pathological nature of the Gordian graph makes it usually difficult to pinpoint its properties. For example, it is still unknown whether the Gordian graph is homogeneous; see Budney [4]. Also, figuring out the path distance between two knots is regarded as a hard problem (the computation of unknotting numbers is a subproblem).

The aim of this paper is to study the opposite point of view: instead of zooming out on the set of all knots, we will describe a way to observe "under the microscope" each knot type.

To this end, we associate to each knot $K \subset S^3$ a graph, the *Reidemeister graph* $\mathcal{G}(K)$, having a vertex for each diagram of K, and an edge between two diagrams whenever one can be converted into the other by a single Reidemeister move. We will make these definitions precise in the next section, but we note here that, unlike the Gordian

Published: 23 April 2020 DOI: 10.2140/agt.2020.20.643

graph, each $\mathcal{G}(K)$ is locally finite, so it is in some sense better behaved, and many of its properties can be studied through combinatorial techniques.

The definitions and proofs in this paper are quite natural and elementary in spirit: the only nontrivial tools we are going to use are the diagram invariants of Arnold [2] and Hass and Nowik [10] introduced in Section 2. An analogous construction in a slightly different context (see Section 4) has been carried out by Miyazawa [15]. Subtle differences in our initial setting will allow us to prove Theorem 1.1 (see also Question 6.6).

The paper is structured as follows: Section 2 gives the basic definitions of the planar and S^2–Reidemeister graphs, collectively called $\mathcal{R}$–graphs.

In Section 3 we analyse some local properties of these graphs; in particular we will classify all short paths in them (Theorem 3.2), and examine the change in valence between adjacent vertices. These technical results are going to be crucial to establish the main result of the paper. As a preliminary step we will prove in Theorem 3.23 that the graph can detect which Reidemeister move corresponds to each of its edges, and define a related notion of *diagram complexity*.

Section 4 instead deals with global properties of the $\mathcal{R}$–graphs; we show that unsurprisingly they are nonplanar (Figure 43) and not hyperbolic (Proposition 4.1). In addition we show that each Reidemeister graph has only one thick end (Proposition 4.3), and compute the homology groups of an associated simplicial complex, following the definition of Miyazawa [15].

In Section 5 we are going to prove the main result of this paper, concerning the completeness of the S^2–Reidemeister graph invariant:

Theorem 1.1 *The S^2–Reidemeister graph is a complete knot invariant up to mirroring; that is, $\mathcal{G}(K) \equiv \mathcal{G}(K')$ if and only if K' is isotopic to K or $\overline{K}$.*

Indeed the proof of this theorem will guarantee a stronger result (Proposition 5.3): the isomorphism type of the graph does not only distinguish all knots, but contains enough information to recover some diagrams of the knot (up to mirroring). Moreover, all this data can be extracted from finite portions of the graph (Corollary 5.4).

We remark that, unlike the previously known complete invariants — such as knot complement in Gordon and Luecke [8], quandles in Joyce [11] and conormal tori in Ekholm, Ng and Shende [6] — the proof of completeness for the S^2–graph is substantially more elementary, and self-contained.

Finally, in Section 6 we define yet another kind of graph, relating the Gordian and Reidemeister graphs by a "blowup" construction.

Acknowledgements The authors would like to thank C Collari, P Lisca, A Juhász, P Ghiggini, M Golla, F Lin, and D Eridan for stimulating conversations. The authors are extremely thankful to M Lackenby for pointing out some mistakes in an early draft and for his continuous help. A special thanks to M Marengon, for his remarks, suggestions and valuable help. Barbensi would like to thank her other supervisors Dorothy Buck and Heather Harrington for their support. Barbensi is supported by the EPSRC grant "Algebraic and topological approaches for genomic data in molecular biology" EP/R005125/1.

Celoria has received support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement no 674978).

The authors would also like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the program "Homology theories in low dimensional topology", supported by EPSRC grant no EP/K032208/1, where work on this paper was undertaken.

2 The graph

We start by giving some precise definitions of the well-known objects we are going to use extensively in the following.

As usual by *knot* we mean the ambient isotopy class of a tame embedding of S^1 in S^3. The set of unoriented knot types in S^3 will be denoted by $\mathcal{K}$, and the set of diagrams representing a knot K by $\mathcal{D}(K)$. A knot diagram $D \in \mathcal{D}(K)$ can be thought of as a 4–valent graph in $\mathbb{R}^2$ or S^2, by disregarding the crossing information. In order to avoid confusion, we are going to refer the 4–valent graph associated to a diagram as the *knot projection*. We will call an *arc* each portion of a diagram or projection which connects two crossing points, and denote by $\alpha(D)$ the number of arcs in D. By the handshaking lemma we have $\alpha(D) = 2\,\mathrm{cr}(D)$, where $\mathrm{cr}(D)$ is the number of crossings in D. From now on, we are going to assume that, unless otherwise stated, each diagram D contains at least one crossing.

The complement of a planar knot projection is composed of polygons, with the exception of the "external" region which is a punctured polygon; we will call this external part a

Figure 1: Example of a local planar isotopy acting on an arc.

polygon as well. As is customary we denote by $p_k(D)$ the number of polygons with k–edges.

We have

$$\alpha(D) = \frac{1}{2}\sum_{k\geq 1} k \cdot p_k(D), \tag{2-1}$$

and the number of regions in $S^2 \setminus D$ is $\sum_{k\geq 1} p_k(D)$.

We say that a planar diagram D is *periodic* if there exists a nontrivial rotation of the projection plane taking D to itself, and a knot K is periodic if it admits a periodic diagram. The *order of periodicity* is then just the order of the rotation acting on the diagram.

A diagram on the 2–sphere is said to be *periodic* if there is a nontrivial, finite-order, orientation-preserving self-diffeomorphism of the sphere that takes it to itself. The $(2n+1, 2)$–torus knots are an example of knots that exhibit a D_{2n+1} periodicity[1] on the 2–sphere and cyclic periodicity of order 2 or d (with $d \mid 2(2n+1)$) on the plane.

Conversely, a knot which does not admit any periodic diagram is said to be *nonperiodic*.

A *planar isotopy* can modify locally a knot diagram by moving slightly an arc as in Figure 1, or by displacing a whole diagram, without creating or removing any crossing.

Note that by considering $\mathbb{R}^2$ rather than S^2 as the ambient space we get a "larger" set of diagrams; the two diagrams of the left trefoil minimising its crossing number shown in Figure 2 are planar isotopic on the 2–sphere but not on the plane.

Probably the most fundamental result in knot theory is Reidemeister's theorem [17],[2] stating that two diagrams represent the same knot type if and only if they are related by a finite sequence of local moves, known as Reidemeister moves, together with planar isotopy. These moves are described in Figure 3. Note that this set of moves is

[1] Here D_m denotes the dihedral group of order $2m$.

[2] We will adhere to the standard attribution of the theorem, which in fact was independently discovered by Alexander and Briggs [1].

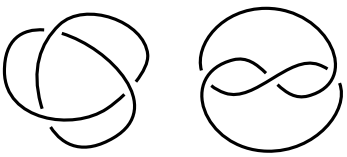

Figure 2: These two diagrams of the left trefoil are planar isotopic on S^2, but not on $\mathbb{R}^2$.

not *minimal* (see [16] for the statement in the oriented case); in fact one of the two Ω_1 moves could be discarded (Figure 10). However the choice of this slightly larger set will be crucial in the proof of all the upcoming results. In what follows we will find it convenient to divide the Ω_2 moves into two kinds. The first kind consists of those Ω_2 moves performed on the configuration in Figure 4, which is called a *tentacle*. We will denote them by Ω_T, and call them *tentacle moves*. In other words, Ω_T moves are the Ω_2 moves that create a tentacle (and their inverses), as in Figure 5. We will say that a tentacle configuration has *height* m if it can be expressed as the composition of $m+1$ Ω_1 moves (with alternating signs). In particular a tentacle of height 1 is the result of performing two Ω_1 moves with opposite curls, or an Ω_T move in which the two affected strands belong to the same arc (Figure 5, top). A tentacle of height m contains $m-1$ subtentacles of heights $m-1,\dots,1$ as subconfigurations.

The other kind (which we will simply call Ω_2) instead is any other Reidemeister move of type 2. The reason for this distinction will become apparent in the next sections (see Theorem 3.2); in fact we are going to prove that tentacle moves are intrinsically distinguished from the other moves (Theorem 3.23).

Additionally, if Ω denotes a Reidemeister move which is not an Ω_3, there are two cases, according to whether we are doing or undoing the move. Hence, when necessary,

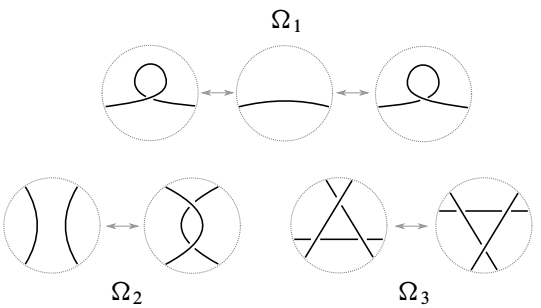

Figure 3: The standard Reidemeister moves on knot diagrams.

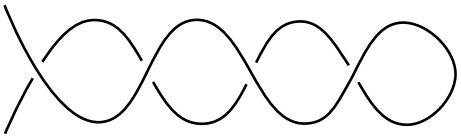

Figure 4: A tentacle configuration of height 3. Note that the crossings have alternating signs.

we are going to denote a move by Ω^+ or Ω^- if it increases (respectively decreases) the crossing number.

Recall that the *writhe* of a diagram D is the sum of the signs of the crossings; it changes by ± 1 when an Ω_1 move is performed, and is left unchanged by the other moves. The last definition we will need is the *mirror* $\overline{D}$ of a diagram D, which is just the diagram obtained by switching all crossings in D. A knot is said to be *amphicheiral* if it is unchanged under mirroring of its diagrams.

We are ready to introduce the two basic versions of the object we are going to study throughout the rest of the paper.

Definition 2.1 Given a knot $K \subset S^3$ define the *Reidemeister graph* of K, denoted by $\mathcal{G}(K)$, as the graph whose vertices are the diagrams of K on the 2–sphere (up to isotopies of S^2) and which has an edge between two diagrams if and only if they are connected by a single Reidemeister move. If we replace diagrams in S^2 with planar diagrams, we obtain the planar Reidemeister graph $\mathcal{G}_P(K)$.

In what follows we will use the term "Reidemeister graph" or $\mathcal{R}$–graph to denote both $\mathcal{G}(K)$ and $\mathcal{G}_P(K)$.

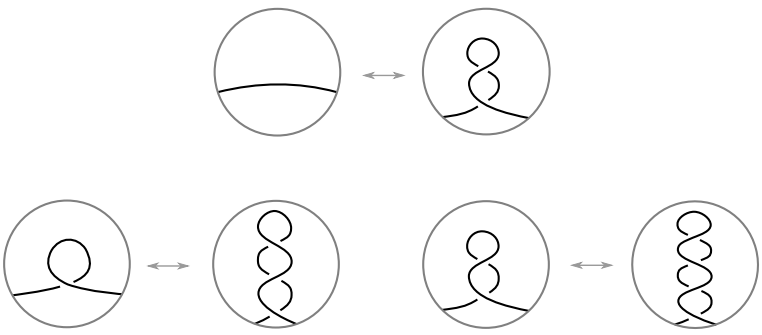

Figure 5: The tentacle moves. In the top part of the figure a tentacle move creating a tentacle configuration of height 1 is shown; such a move always arises as the superposition of an arc on itself.

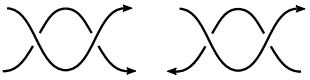

Figure 6: A matched Ω_2 move (left) and an unmatched Ω_2 move (right).

Remark 2.2 It might happen that two diagrams are connected by two different moves (see eg Figure 8), which will be considered as different edges in the graphs except that moves coinciding up to a planar isotopy will be represented by a single edge.

There are a few immediate consequences of this definition; first of all the isomorphism class of the graphs $\mathcal{G}(K)$ and $\mathcal{G}_P(K)$ are knot invariants, and they are unchanged under mirroring of the knot.

Also, Reidemeister's theorem implies that for each $K \in \mathcal{K}$ the corresponding $\mathcal{R}$–graphs are connected.

It might seem strange to define invariants that are more complicated than the object we started with. However, beyond their intrinsic interest, the $\mathcal{R}$–graphs will allow us to produce several related simple numerical invariants.

To prove many of the local structure results of Section 3 for the graphs $\mathcal{G}(K)$ and $\mathcal{G}_P(K)$, we will need the diagram invariant introduced by Hass and Nowik in [10], whose properties are concisely recalled below. In the specialised form we are going to use it, this invariant takes values in the free abelian group generated by the formal variables $\{X_s, Y_s\}_{s\in\mathbb{Z}}$. We call an Ω_1^+ positive if the crossing (for an arbitrary choice of orientation) is positive, and negative otherwise; we will call an Ω_2 move *matched* if the two strands go in the same direction, and *unmatched* otherwise (see Figure 6). Note that Ω_T moves are always unmatched.

If two diagrams D and D' differ by a single Reidemeister move, the corresponding I_{lk}–invariants differ as shown below, for some $n, m, r \in \mathbb{Z}$:

$$I_{\mathrm{lk}}(D') = \begin{cases} I_{\mathrm{lk}}(D) + X_0 & \text{if the move is a positive } \Omega_1^+, \\ I_{\mathrm{lk}}(D) + Y_0 & \text{if the move is a negative } \Omega_1^+, \\ I_{\mathrm{lk}}(D) + X_n + Y_{n+1} & \text{if the move is a matched } \Omega_2^+, \\ I_{\mathrm{lk}}(D) + X_m + Y_m & \text{if the move is a unmatched } \Omega_2^+, \\ I_{\mathrm{lk}}(D) + X_0 + Y_0 & \text{if the move is an } \Omega_T^+, \\ I_{\mathrm{lk}}(D) + \begin{cases} \pm(X_r - X_{r+1}) \\ \pm(Y_n - Y_{n+1}) \end{cases} & \text{if the move is an } \Omega_3. \end{cases}$$

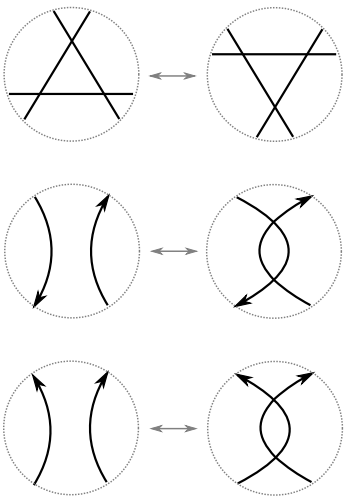

Figure 7: A triple-point perestroika, followed by the two possible self-tangency perestroikas.

The only other nontrivial invariants we are going to use are Arnold's *perestroika* invariants St and $J^{\pm}$, first defined in [2]. These are invariants of regular homotopy classes of immersions of S^1 in $\mathbb{R}^2$ or S^2. They change in a controlled way under perestroikas, that is, the analogues of Reidemeister moves for immersions $S^1 \looparrowright \mathbb{R}^2$ (or S^2), as shown in Figure 7. We will use them on the knot projections associated to the diagrams. Note that there is no analogue of Ω_1 moves for immersions, since performing it would change the *index* of the curve.

The invariant St changes by ± 1 under a triple-point perestroika (which corresponds to an Ω_3 move in our setting), and is left unchanged under self-tangency perestroikas (corresponding to Ω_2 moves and Ω_T moves). On the other hand, the invariant J^+ is unchanged under triple-point perestroikas and changes by a fixed positive amount (conventionally 2), when a direct-tangency perestroika is performed (that is, a matched Ω_2^+). The invariant J^- behaves similarly, but changes only for inverse self-tangency perestroikas (that is, an unmatched Ω_2^+ or Ω_T^+ in our case).

3 Local properties

Given a knot $K \in \mathcal{K}$ and $D \in \mathcal{D}(K)$, the Reidemeister graphs can be naturally endowed with the path metric. Note that the distance induced by this metric coincides with the minimal number of Reidemeister moves connecting two diagrams. We denote

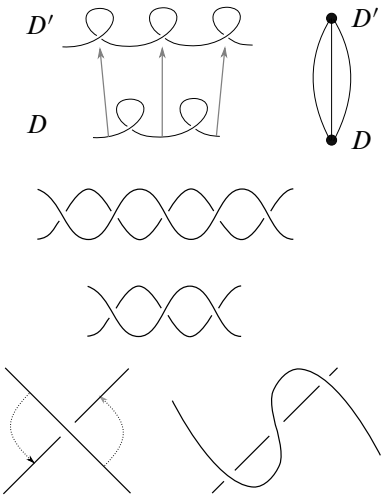

Figure 8: In the top part some curls on an arc inducing a multiedge on the graph, together with the corresponding configuration. In the central part, a multiedge induced by Ω_2 (or Ω_T) moves, and in the lower part a 2 multiedge induced by Ω_2 moves.

by $S(D)$ the subgraph induced by the vertices having distance ≤ 1 from D, and more generally $S_R(D)$ will denote the subgraph spanned by the diagrams with distance $\leq R$ from D. As we will see in what follows, a lot of information about a diagram D can be extracted from $S(D)$.

The next results are aimed at understanding in detail the structure of small portions of the Reidemeister graph, in both the periodic and nonperiodic cases.

We will find it convenient to denote by $\#\Omega_i^{\pm}(D)$ the number of Reidemeister i moves of type $\pm$ which can be applied to D.

This next result states that there are no "cosmetic" Reidemeister moves, meaning that a Reidemeister move necessarily changes the diagram, even up to planar isotopy.

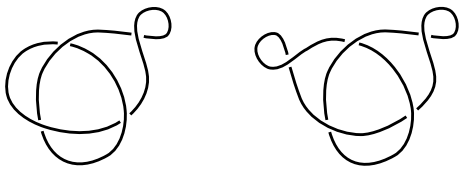

Figure 9: There are two inequivalent Ω_1 moves that take one diagram to the other. Note that, even if the diagrams are not periodic, they represent a periodic knot type.

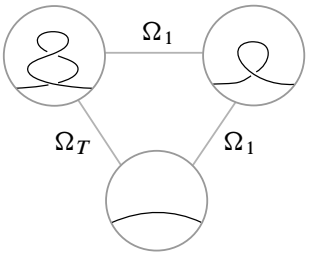

Figure 10: A triangle $\Omega_T^{\pm}-\Omega_1^{\mp}-\Omega_1^{\mp}$.

Proposition 3.1 *The graphs $\mathcal{G}_P(K)$ and $\mathcal{G}(K)$ do not contain any self-edges.*

Proof Since Ω_1 and Ω_2 moves change the crossing number, they can be immediately ruled out. The only possibility is then to have an Ω_3 move that, if performed, takes a diagram $D \in \mathcal{D}(K)$ to itself (up to planar isotopy). It is, however, easy to exclude this case as well using the Hass–Nowik I_{lk} invariant (or Arnold's St): as recalled in the previous section this invariant changes in a nontrivial manner under Ω_3 moves. □

It is easy to realise that for a given knot, its $\mathcal{R}$–graph contains infinitely many multiedges, of any order: just take n identical curls on the same arc for one diagram and $n+1$ on the other. Then there are $n+1$ edges connecting them, corresponding to the possible choices for adding another curl, as shown in the top part of Figure 8 for $n = 2$.

It is also possible to find multiedges induced by Ω_2 moves, as shown in the middle and lower parts of Figure 8.

In fact, using the configuration in the lower part of Figure 8, it is immediate to show that the only radius 1 ball not containing multiedges is centred in the crossingless diagram of the unknot.

If the knot is periodic, one can also have multiedges of the form shown in Figure 9. It is, however, easy to prove[3] that each multiedge must be composed of moves of the same kind.

We will say that a graph contains a *triangle* if there are three distinct vertices such that each vertex is at distance 1 from the other two.

We want to analyse the shape of the cycles in $S(D)$, other than the multiedges. It is easy to find a cycle of length 3, shown in Figure 10. Moreover, since this cycle can

[3]For example, using Arnold's invariants for Ω_2 and Ω_3 moves, and I_{lk} for Ω_1 moves.

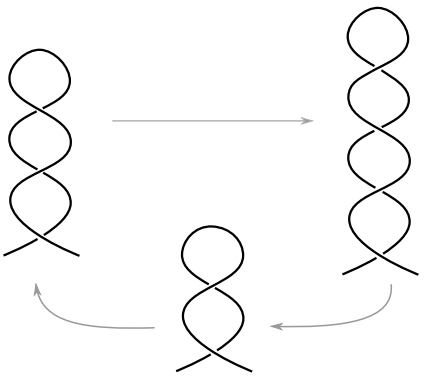

Figure 11: A triangle involving some tentacle configurations.

start from any unknotted portion of an arc, it is ubiquitous in all Reidemeister graphs. A similar and slightly more elaborate example involving a "higher" tentacle is shown in Figure 11. The following result will establish that in some sense these are the only possible cases. Moreover it will permit us to explore the main properties of the graph. Its proof is roughly based on the following idea: the total sum of any diagram invariant has to vanish on a closed cycle. In most cases it will be sufficient to consider very simple diagram invariants, such as the crossing number.

Theorem 3.2 *If K is a nontrivial knot, the only triangles in its Reidemeister graphs are of the form $\Omega_T^{\pm}-\Omega_1^{\mp}-\Omega_1^{\mp}$. If instead K is the unknot $\bigcirc$, there are some sporadic exceptions* (shown in Figure 12) *of cycles of the form $\Omega_2^{\pm}-\Omega_1^{\mp}-\Omega_1^{\mp}$.*

Proof Suppose we have a length 3 cycle, connecting the diagrams D_0, D_1 and D_2. The total change of crossing number must be 0, hence we can immediately exclude

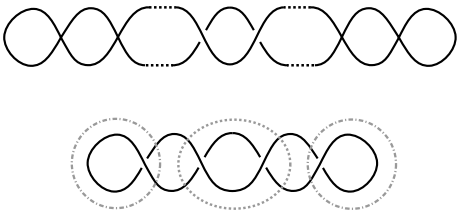

Figure 12: On the top, a diagram for the unknot obtained from the crossingless one by an $\Omega_{2/T}$ move followed by successive Ω_1^+ moves creating crossings of any sign. On the bottom, an example of a triangle involving diagrams of this kind: performing the central Ω_2^- move or the two lateral Ω_1^- moves produces the same diagram.

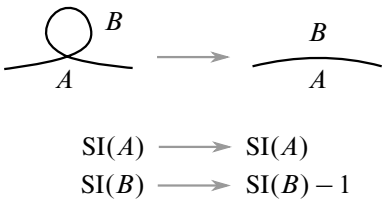

Figure 13: Undoing an Ω_1 move decreases $\mathrm{SI}(B)$ by 1 and leaves $\mathrm{SI}(A)$ unchanged.

most cases; a priori the only possible combinations (up to permutations) of three Reidemeister moves that could work are

(1) Ω_3–Ω_3–Ω_3,

(2) Ω_3–Ω_2–Ω_2,

(3) Ω_3–Ω_T–Ω_2,

(4) Ω_3–Ω_T–Ω_T,

(5) Ω_3–Ω_1–Ω_1,

(6) Ω_2–Ω_1–Ω_1,

(7) Ω_T–Ω_1–Ω_1.

It is easy to exclude cases (1) to (4) using Arnold's St invariant: in any cycle (not containing Ω_1 moves) the number of Ω_3 moves must be even. Case (5) can instead be excluded using the Hass–Nowik invariant: the Ω_3 move contributes to I_{lk} with two consecutive terms (that is, terms of the form $A_n - A_{n+1}$ for $A = X$ or Y and $n \in \mathbb{Z}$), while the Ω_1 moves can only add some terms of the form $\pm A_0$. Hence the total change in the sum cannot be 0.

Finally we can focus on cases (6) and (7) and exclude the former. First notice that, in order to preserve the crossing number, an $\Omega_2^{\pm}$, must be followed by two $\Omega_1^{\mp}$. Moreover, using the Hass–Nowik invariant we can conclude that the crossings involved in the Ω_1 moves have different signs, and that the Ω_2 is unmatched.

Define the *self-intersection number* $\mathrm{SI}(P)$ of a region P in the complement $S^2 \setminus D$ as the number of crossings in the boundary of P that connect P to itself. We can associate to each diagram D an unordered N–tuple $\mathrm{SI}(D) = (\mathrm{SI}(P_1), \mathrm{SI}(P_2), \dots, \mathrm{SI}(P_N))$ where N is the number of regions in $S^2 \setminus D$.

Performing an Ω_1^- move always decreases the self-intersection number of a single region by 1, and leaves the self-intersection numbers of the other regions unchanged; see Figure 13.

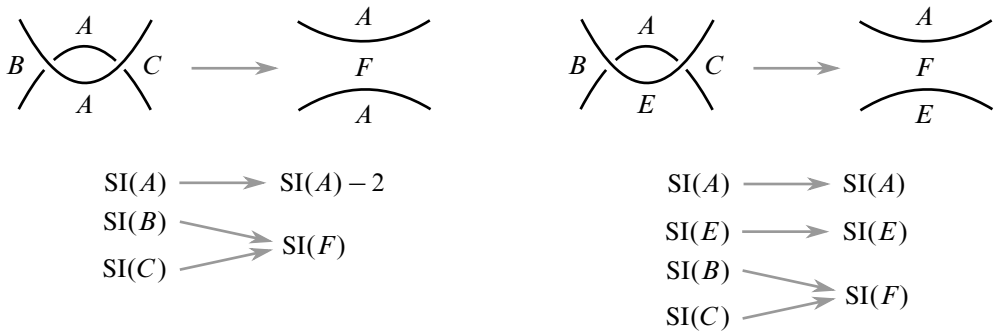

Figure 14: Changes in the self-intersection numbers for an Ω_2^- move.

On the other hand, an Ω_2^- move can change the self-intersection N–tuple in two different ways, depending on whether the regions denoted by A and E in Figure 14, right coincide or not.[4] If A and E coincide, then the component $\mathrm{SI}(A)$ of $\mathrm{SI}(D)$ decreases by 2 when the move is performed (as in Figure 14, left). In the other scenario, the only change in $\mathrm{SI}(D)$ comes from the merging of the regions B and C; the new region formed has as self-intersection number greater than or equal to $\mathrm{SI}(B)+\mathrm{SI}(C)$ (Figure 14, right).

Suppose now, for a contradiction, that there exists a cycle of the form $\Omega_2^{\pm}-\Omega_1^{\mp}-\Omega_1^{\mp}$. That means we can obtain a diagram D' from D either by performing an Ω_2^- move or a sequence of two Ω_1^- moves on D and that the changes in $\mathrm{SI}(D)$ must be the same. Now observe that, while the self-intersection number of at least one region decreases with two consecutive Ω_1^- moves, if in the Ω_2 move the regions A and E are distinct, the sum of the self-intersections over all regions is increased or left unchanged. This fact allows us to exclude the case in which the Ω_2 move is as in Figure 14, right.

We will find it useful to divide the discussion in cases, depending on the mutual positions of the curls undone by the Ω_1^- moves. The relevant portions of the initial diagram D_0 are displayed in Figure 15 for each of these possibilities: in the first and second rows we show the mutual positions the curls can have if they do not both appear in D_0; in other words, the 1–region undone by the second Ω_1 appears after undoing the first curl.[5]

In the third row letters indicate the regions touched by the curls: these regions can either coincide or not.

[4]The regions denoted by B and C are always distinct, otherwise the diagram would represent a two-component link.

[5]Recall that these two moves must have opposite signs.

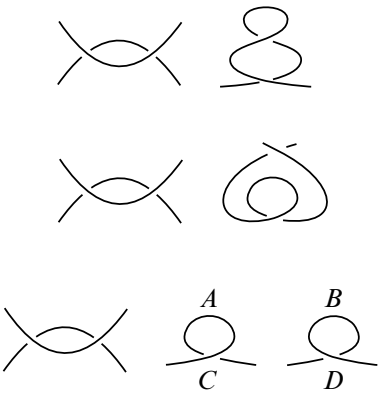

Figure 15: In each row we show the portions of D_0 involved in the Ω_* moves. In the first and second rows, which refer to as cases (1) and (2), we assume that the curls undone by the Ω_1 moves do not both appear in the diagram, while in the third row, case (3), they do. Letters in the third row indicate the regions touched by the curls.

In what follows, for each one of these cases we will either prove that the Ω_2^- needs to be a tentacle move (Figure 5), or exclude the configuration.

The first case in Figure 15 can be settled as follows: Consider Figure 16. We can see that in the diagram D_2 there is a tentacle appearing. Since the diagrams D_1 and D_2 are equivalent by hypothesis, the tentacle in D_2 must appear somewhere in D_1. Moreover, since they coincide out of the portions of diagram drawn in the figure, the presence of a tentacle in D_1 *far away*[6] from the portions drawn would imply the existence of an identical tentacle somewhere in D_2, and we would still have one more tentacle in D_2 than in D_1.

A similar recursive argument applies if the tentacle appears by undoing the Ω_1 moves, as in Figure 17. In fact, in each of the cases shown in Figure 17, there is a configuration in D_2 which does not appear in D_1, and the only way to have $D_1 = D_2$ is to find this configuration in D_1. Iterating this procedure, one sees that the two diagrams cannot be equivalent.

It follows that the only way D_1 can be equivalent to D_2 is if the Ω_2 is in fact an Ω_T; thus the corresponding part of the diagram is a portion of a tentacle.[7]

[6]Here and in what follows, by "far away" we mean that the configuration is left untouched by the moves considered.

[7]If the portion of diagram involved in the Ω_2^- is attached to a piece whose projection is the same as a tentacle, but with "wrong" crossings, then the diagram does not fit in a triangle.

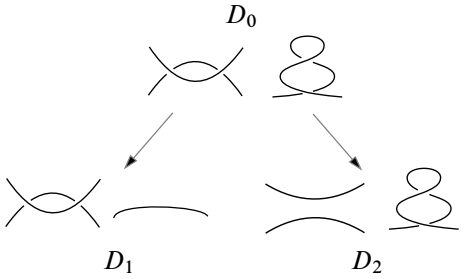

Figure 16: D_1 is the diagram obtained after performing the two Ω_1^- moves: together they cancel the tentacle appearing in D_0. D_2 is the result of undoing the Ω_2 move in the left-hand portion of the diagram in D_0.

For the second case consider Figure 18: we apply the same argumentation of case (1). Since the diagrams D_1 and D_2 are equivalent the "heart-shaped" configuration in D_2 must appear somewhere in D_1. Moreover, since the diagrams coincide out of the portions drawn in the figure, the presence of a heart in D_1 *far away* from the portions drawn would imply the existence of an identical heart somewhere in D_2, and we would still have one more heart in D_2 than in D_1. The same argument of case (1) (as in Figure 17) works if we assume that the heart appears after undoing two Ω_1 moves. It follows that the only possibility is the one depicted in Figure 19.

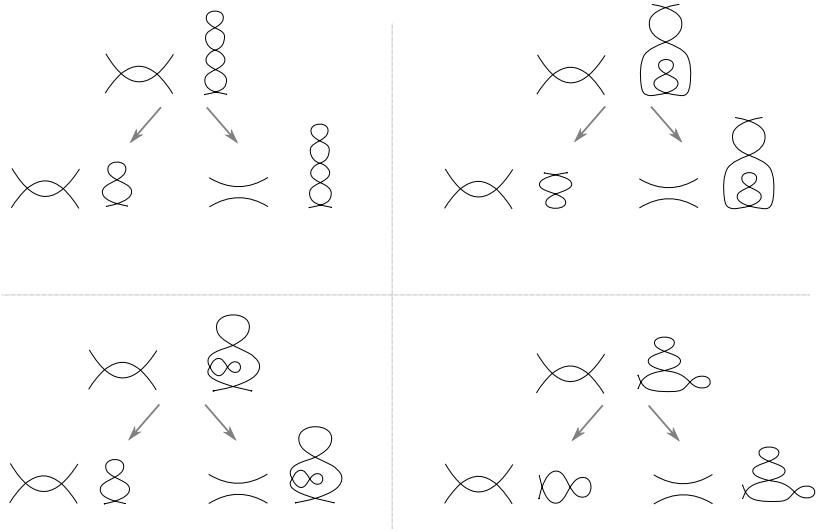

Figure 17: The only possible ways a tentacle can appear in case (1) after performing two Ω_1^- moves. In each of these cases we can exclude that the diagrams form a triangle by a recursive argument.

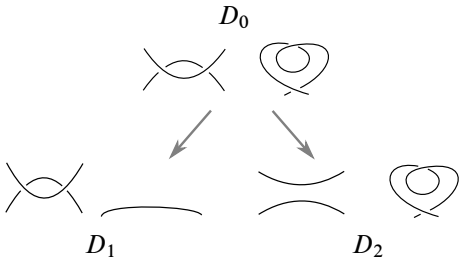

Figure 18: D_1 is the diagram we obtain after performing the two Ω_1^- moves: together they cancel the heart configuration appearing in D_0. D_2 is the result of undoing the Ω_2 move in the left-hand portion of diagram in D_0.

We can, however, prove that in this case D_1 and D_2 cannot be equivalent diagrams, and thus exclude it. To this end, consider the blackboard framing of the projection: there are two possibilities to be considered, since we can draw the framing curve on either side of the diagram. Then, since we do not know how the portions of diagrams involving the moves are positioned with respect to each other, we need to consider four different cases, all shown in Figure 20. It is easy to argue that D_1 and D_2 cannot be equivalent, since the number of curls having the blackboard framing "inside" is different in all four cases.

We are now left with case (3) from Figure 15. As usual, it is convenient to have in mind all the diagrams involved in the triangle, as in Figure 21.

From Figure 21 it is apparent that there are two more visible 1–regions in D_2 than in D_1: since by hypothesis the diagrams are equivalent, there must be two curls in D_1 as well. Suppose, for a contradiction, that the Ω_2 is not a tentacle (that is, $b, c \neq 1$).

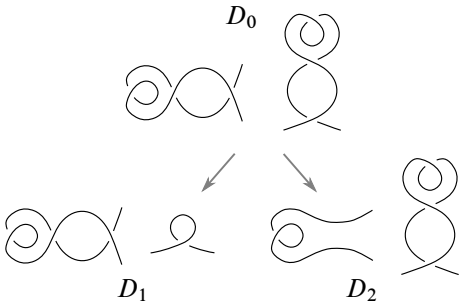

Figure 19: D_1 is the diagram we obtain after performing the two Ω_1^- moves: together they cancel the heart configuration appearing in D_0, leaving a curl. D_2 is the result of undoing the Ω_2 move in the left-hand portion of diagram in D_0.

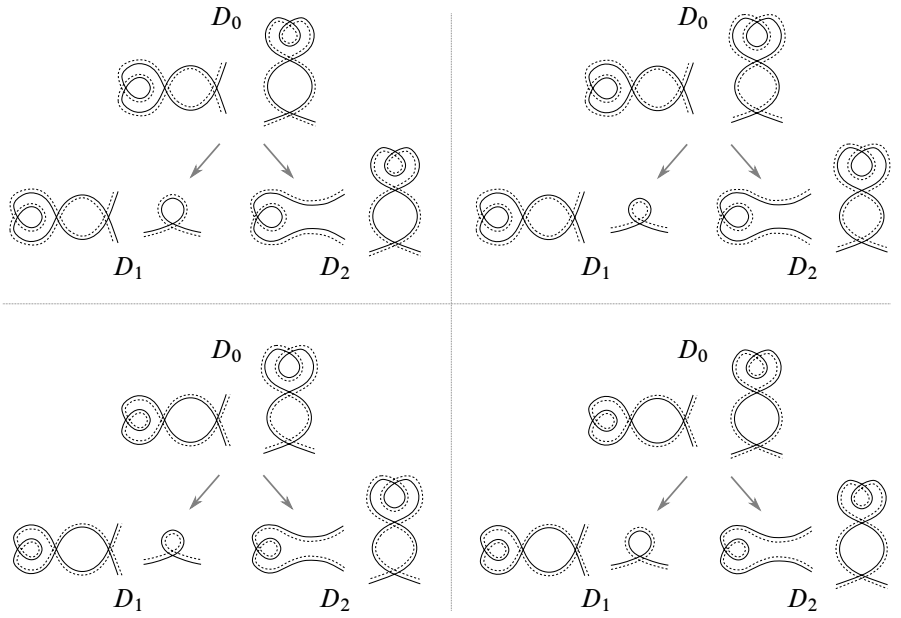

Figure 20: The four possible choices for the blackboard framing.

Then, the straight lines in D_1 left by undoing the Ω_1 moves must be part of two curls. If we assume (Figure 21) that the regions touching the curls in D_0 are different, this means that at least two among $d-1$, $e-2$, $f-1$ and $g-2$ must be equal to 1. Since the cases $(e,d)=(3,2)$ and $(f,g)=(2,3)$ are impossible, we are in one of the cases described in Figure 22.

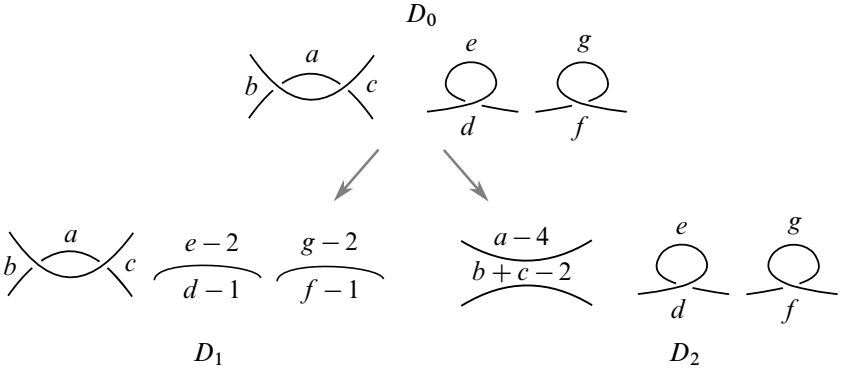

Figure 21: Lowercase letters indicate the number of edges in each region. Keep in mind that, even if in the picture all the regions are depicted as different, some of them might coincide.

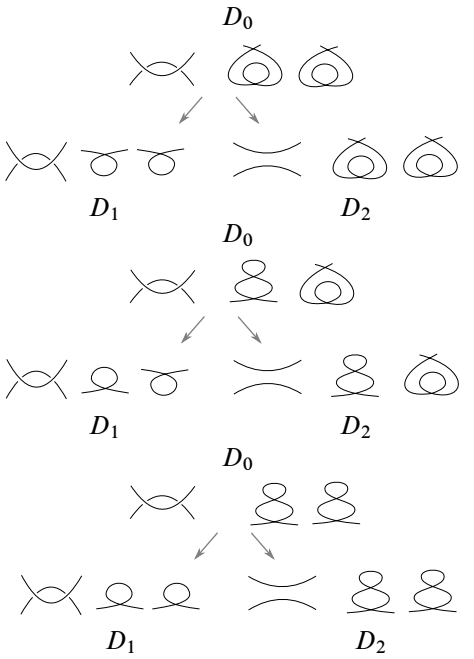

Figure 22: The three possible kinds of triangles, assuming that the regions touching the curls undone by the Ω_1 moves do not coincide.

Before dealing with the configurations described in Figure 22, we need to consider the cases in which some of the regions touching the curls coincide, keeping in mind that we are assuming that the Ω_2 is not a tentacle move. We have the following possibilities (capital letters denote regions, as in Figure 15):

(I) $A = B$ and $C = D$;

(II) $A = D$ and $B = C$;

(III) $A = B$ and $C \neq D$;

(IV) $C = D$ and $A \neq B$;

(V) $A = D$ and $B \neq C$;

(VI) $C = B$ and $A \neq D$.

Note that the upper and lower regions left by undoing a curl cannot be both 1–regions. Thus, cases (I) and (II) are straightforward to exclude, since if the regions coincided pairwise it would be impossible to recover two 1–regions from the straight lines left by undoing the starting curls. For the same reason, in the third case the only way to have two curls left after the Ω_1^- moves have been performed is to have a 2–region

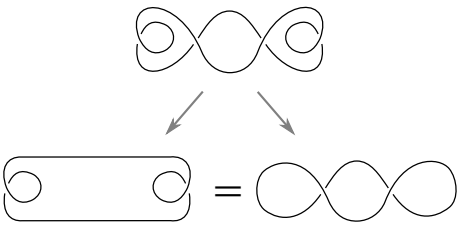

Figure 23: A triangle for the unknot fitting in the family described in the statement of Theorem 3.2.

below each 1–region in D_0. Thus, case (III) fits in the bottom configuration described in Figure 22. Similarly, in case (IV), we would necessarily have both the curls in D_0 lying inside a 4–region, forming a *heart* and fitting in the top case shown in Figure 22. The latter two cases are symmetric, and it is enough to discuss only the first one. Again, since it is impossible to have both the upper and lower region left by undoing a curl as 1–regions, it follows that the Ω_1^- moves must be performed in portions of diagrams identical to the ones drawn in the middle case of Figure 22.

Let's now discuss carefully Figure 22. Consider the configuration at the top of the figure: since the diagrams D_1 and D_2 are equivalent, the *heart* configurations in D_2 must appear somewhere in D_1. Moreover, since they coincide out of the portions drawn in the figure, and since using again a recursive argument we can exclude that they are created by undoing the curls in D_0, the only possibility is that these hearts are attached to the Ω_2–portion in D_1, as shown in Figure 23.

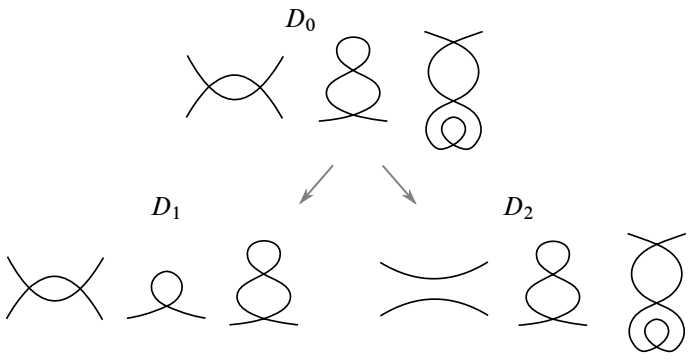

Figure 24: D_1 is the diagram we obtain after performing the two Ω_1^- moves: with the first one we cancel the curl inside the heart, while the other has the effect of decreasing the height of the left tentacle by 1. D_2 is the result of undoing the Ω_2 move in the left-hand portion of D_0.

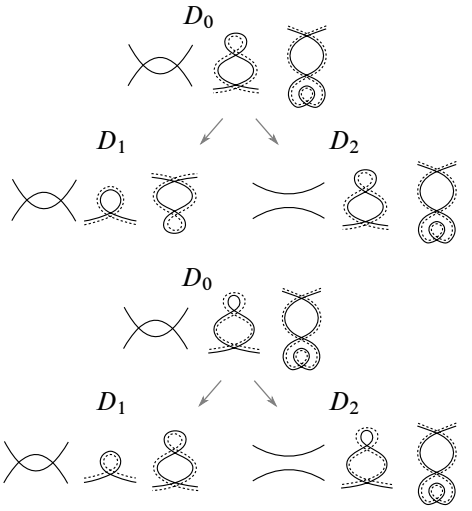

Figure 25: Two of the four possible choices of framing. The remaining two can be treated in the exact same way.

Note that even if in this case D_1 and D_2 turn out to be equivalent, the diagrams represent the trivial knot, and more precisely they fit in the family described in the statement of the theorem.

Notice that this can only happen if we are working with diagrams on S^2; if we are working with planar diagrams instead, this configuration does not fit in a triangle. Now, call *generalised tentacles* the configurations formed by two successive Ω_1 moves made one on top of the other, as appearing in D_2 and D_0 on the bottom of Figure 22. If the crossings are such that the configurations form tentacles, then this implies (as in case (1) of Figure 15) that the Ω_2 is in fact a tentacle move.

Otherwise, by using a similar recursive argument as before, together with the fact that the upper and lower regions involved in the Ω_2 move coincide, we can exclude both the possibility that the configurations appear in D_1 by performing the Ω_1^- moves, and that they appear somewhere *far away* from the portions of diagram shown. Thus, we see that the only possibility for D_1 and D_2 to be equivalent occurs when the generalised tentacles are attached[8] to the Ω_2–portion of D_2, forming a diagram for the unknot fitting in the family described in the statement of the theorem (see Figure 12). Notice that the triangle in $\mathcal{G}(\bigcirc)$ involving the heart configurations described before is a special case of this situation.

[8]Or they are part of longer generalised tentacles attached to the Ω_2–portion of D_1.

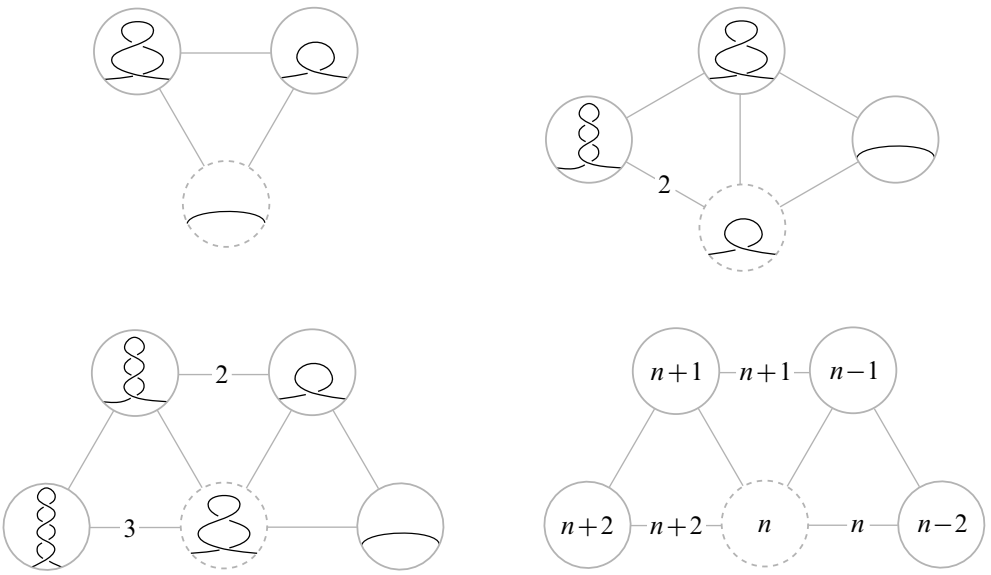

Figure 26: Some of the possible configurations in $S(D)$ involving at least one Ω_1 move (D is contained in the dotted circles). The top left one is present in any $S(D)$, while the others can be found whenever there is an Ω_1^- (top right), a height 1 tentacle (bottom left), or a tentacle of height $n \geq 2$ (bottom right). Numbered edges denote the valence of the corresponding multiedge.

Finally, we are left with the middle configuration in Figure 22. As usual, since D_1 and D_2 are equivalent by hypothesis, the tentacle configuration in D_1 has to appear somewhere in D_2 as well. Assuming that the Ω_2 is not a tentacle move, since the diagrams coincide out of the portions drawn, using yet again a recursive argument we can exclude that the tentacle is created by undoing the curls in D_0; hence the only way to have a tentacle in D_1 is the one shown in Figure 24. We can, however, exclude this case as well by adding the blackboard framings. In Figure 25 two of the possible choices of framings are displayed: in both cases D_1 and D_2 are nonequivalent diagrams, since the framings do not coincide on the tentacles or in the 1–regions left. □

Remark 3.3 In what follows, unless otherwise specified, all the results will hold for every knot type with the exception of the unknot ◯.

For each diagram $D \in \mathcal{D}(K)$, the subgraph $S(D)$ consists of triangles (possibly attached to one another) with one vertex in D, and edges emanating from D. Each of these might be a multiedge. If we want to study the possible configurations in $S(D)$

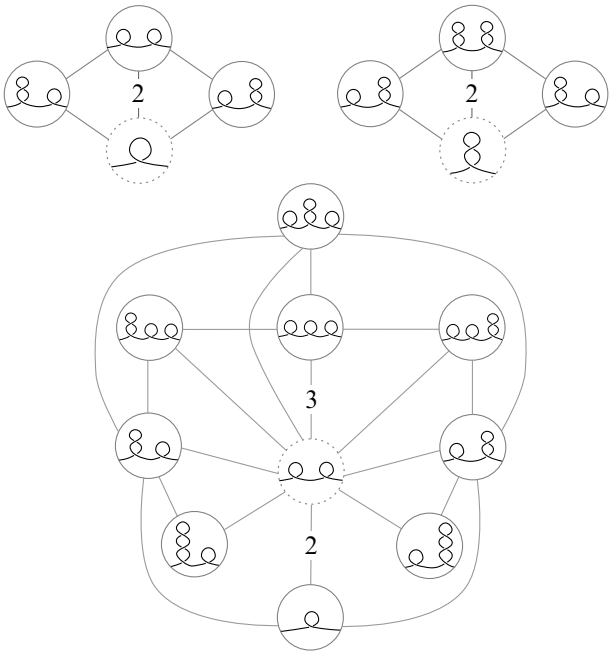

Figure 27: Other qualitatively different triangle configurations formed by $\Omega_T - \Omega_1 - \Omega_1$ can be found whenever there are multiple curls or height 1 tentacles on the same arc.

involving triangles, by Theorem 3.2, we just need to restrict to those containing at least one Ω_1 move; various possibilities involving one or more curls/tentacles are shown in Figures 26 and 27.

So we have a complete description of the short paths that can appear in $\mathcal{G}(K)$; note that it makes less sense to pursue a systematic study of longer (≥ 3) cycles, since any pair of "distant" moves on a diagram produces a cycle of length 4. In the following we are going to examine more closely the properties and shapes of the various triangles that have been produced during the proof of Theorem 3.2. This technical analysis is going to be crucial in the proof of the results leading to Theorem 1.1.

Definition 3.4 We will call a triangle *normal* if it is of the form described in Figure 10, meaning that all the Reidemeister moves are performed locally on the same arc.

Lemma 3.5 *$p_1(D) = 0$ if and only if all the triangles in $S(D)$ are normal.*

Proof If $p_1(D) \neq 0$, then there are at least two triangles sharing an Ω_1^+ edge, as shown in the top-right part of Figure 26. This implies that there are at least two

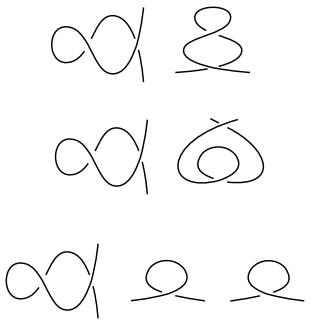

Figure 28: In each row the portions of D_0 involved in the Ω_* moves are shown. In the first and second rows, which we refer to as cases (1) and (2), we assume that the curls undone by the Ω_1 moves do not both appear in the diagram. Case (3) is the third row.

nonnormal triangles, since one can perform the first Ω_1^+ move on either side of the preexisting twirl, and complete this edge to a triangle by performing the successive Ω_1^+ and Ω_T^- on the preexisting twirl.

To show the converse, suppose that $p_1(D) = 0$. Thanks to Theorem 3.2 we know that all triangles are of the form $\Omega_T - \Omega_1 - \Omega_1$; moreover, every Ω_1 and Ω_T is part of at least one triangle. We wish to understand all the possible configurations forming a triangle. To this end, we can use Figure 15, substituting[9] with Ω_T configurations the Ω_2 moves, as in Figure 28.

Since $p_1(D) = 0$, we can exclude the occurrence of cases (2) and (3) of Figure 28. In fact, in each of these triangles, the diagram with lower crossing number admits at least one 1–region. Let's suppose that there exists a nonnormal triangle fitting in case (1) of Figure 28. By definition, this means that the moves are not performed on the same arc. Then, in the lower crossing number diagram, there is at least one 1–region (see Figure 29), contradicting the hypothesis that $p_1(D) = 0$. □

Remark 3.6 If $p_1(D) \neq 0$, then more complicated triangles appear. We show an example of a nonnormal triangle fitting in case (1) of Figure 15 in Figure 30.

In what follows we are going to analyse what happens in the remaining cases. In fact, case (2) of Figure 15 can be excluded as in the proof of Theorem 3.2.

It is convenient to divide the investigation on triangles fitting in case (3) of Figure 15 in two subcases (denoted by (3a) and (3b) respectively), differing in whether or not

[9] We can assume that the Ω_T move happens on the top of the tentacle.

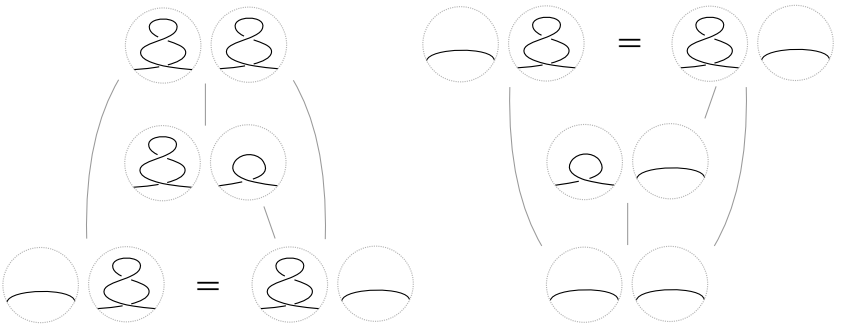

Figure 29: On the left, a nonnormal triangle fitting in case (1). The diagrams with the lowest crossing number (on the bottom) are identified. On the right, a normal triangle fitting in case (1). Here the diagrams with the greatest crossing number are identified, and there is an Ω_T multiedge.

one of the Ω_1^- moves happens on the top part of the tentacle undone by the Ω_T. If it does, then we are in the situation described in Figure 31, and we notice that the diagram with the lowest crossing number contains at least one 1–region; an example of a nonnormal triangle fitting in case (3a) is shown in Figure 32.

Finally, if both the curls undone by the Ω_1 moves are not the top part of the tentacle, then the diagrams appear as in Figure 33.

Again, we can conclude that the diagram with the lowest crossing number presents a tentacle configuration. We show an example of a nonnormal triangle fitting in case (3b) in Figure 34. In all the nonnormal cases above two diagrams are identified, and this implies either the existence of a periodicity for the knot, or that the moves happen on the same edge, involving adjacent curls or tentacles.

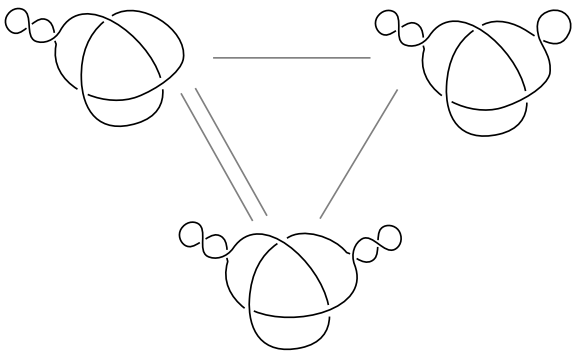

Figure 30: A triangle for a periodic knot fitting in case (1). The left and lower vertices are connected by a multiedge.

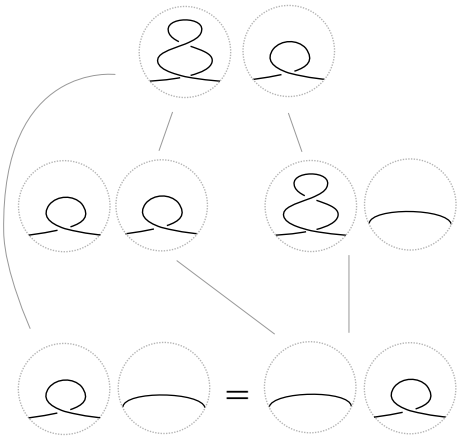

Figure 31: A nonnormal triangle fitting in case (3a). Notice that the diagrams with the lowest crossing number (on the bottom) are identified and present at least a curl.

Lemma 3.7 *Given any knot diagram D, there exists an arc on D such that after performing either an Ω_1^+ or an Ω_T^+ belonging to a normal triangle, the resulting diagrams are nonperiodic. Moreover, if we perform another Ω_T^+ on the top of the tentacle created, the diagram obtained and all of the diagrams in its radius 1 ball are nonperiodic.*

Proof This follows from the fact that the height h tentacle configurations are permuted by any symmetry of the diagram, so if there's only one the diagram cannot be periodic. So, just take any diagram D; if $p_1(D) = 0$, then performing any Ω_1^+ or the Ω_T^+ it is paired with will produce a nonperiodic diagram. If instead D contains at least one curl, choose the one which appears on the top of the highest tentacle, and perform there

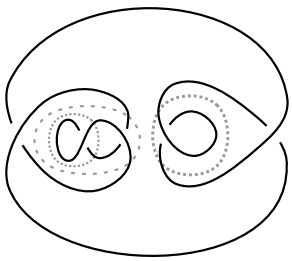

Figure 32: A triangle for a periodic (un)knot, fitting in case (3a). Dotted circles enclose the Ω_1^- moves, and the dashed one the Ω_T^-. This specific example was pointed out by M Marengon.

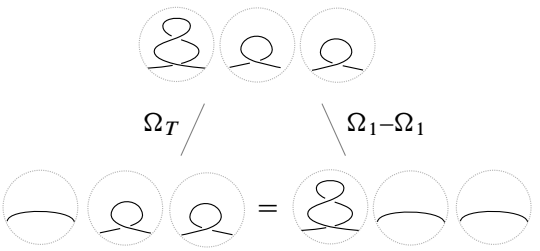

Figure 33: A nonnormal triangle fitting in case (3b). Notice that the diagrams with the lowest crossing number (on the bottom) are identified and present at least a tentacle configuration. See also Figure 34 for an explicit example.

the Ω_1^+/Ω_T^+ pair (with appropriate signs). Since the new diagram will have only one tentacle of maximal height we can conclude. Finally, if we further increase the length of the tentacle, we are sure that we are at least at distance 2 from any periodic diagram. □

Unlike the Gordian graph, the Reidemeister graphs are locally finite, even though the valence is not uniformly bounded (Remark 3.9). The first invariant we can extract from them is in some sense a measure of the minimal complexity of the diagrams of K:

Definition 3.8 Let $v(D)$ denote the valence of the vertex D. The *diagram complexity* of a knot K is

$$\delta(K) = \min_{D \in \mathcal{D}(K)} v(D).$$

If $v(D) = \delta(K)$ we say that D is a *minimum*. We also define $\#\delta(K)$ as the number of minima of $\mathcal{G}(K)$; if a knot type K is such that $\#\delta(K) = 1$, we call K *simple*. Both δ and $\#\delta$ are $\mathbb{N}$–valued knot invariants. There is of course an identical definition for $\mathcal{G}_P(K)$; we denote by δ_P and $\#\delta_P$ the corresponding invariants.

Remark 3.9 We will postpone the proof that $\#\delta(K)$ is in fact well defined until Corollary 3.13. Note that the diagram complexity is not a function of the crossing number, as one might naively think. In Remark 3.22 we are going to provide some examples of this phenomenon. It is, however, true that for a fixed knot type K, the valence becomes arbitrarily high as the crossing number of the diagrams representing K increases.

Given a nonperiodic diagram $D \in \mathcal{D}(K)$, one can enumerate the possible Reidemeister moves on D, in order to compute

$$v(D) = \#\Omega_1^{\pm}(D) + \#\Omega_T^{\pm}(D) + \#\Omega_2^{\pm}(D) + \#\Omega_3(D).$$

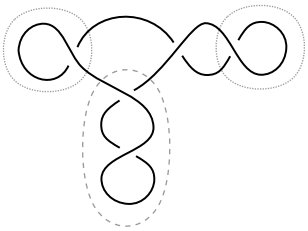

Figure 34: A triangle for a periodic (un)knot, fitting in case (3b).

We start by counting the possible number of Ω_1^+ and Ω_T^+. For each arc in D we can perform four Ω_1^+ moves, as shown in Figure 35, and the same holds for Ω_T^+.

When working with $\mathcal{G}_P(K)$, so with diagrams on the plane, we must put a bit of care in counting Ω_2^+ moves, since the number of such possible moves depends on whether we are in the "external" polygon or not. If a polygon $P \in \mathbb{R}^2 \setminus D$ has k edges, there are $2\binom{k}{2} = k(k-1)$ possible[10] Ω_2^+ moves we can perform in it (the factor of 2 comes from the two possible choices of which arc passes over the other). In the external zone, however, we need to double the previous quantity, since there are two cases to be considered, as shown in Figure 36. So if we denote by k_{ext} the number of edges of the external zone, we have an extra contribution of $k_{\text{ext}}(k_{\text{ext}}-1)$. This extra term does not appear when working with diagrams on the 2–sphere, as there is no preferential polygon.

Adding everything up, we end with this rather unpleasant equation for the valence of a nonperiodic planar diagram:

$$v(D) = 8\alpha(D) + \sum_{k\geq 2} p_k(D)k(k-1) + k_{\text{ext}}(k_{\text{ext}}-1) + \#\Omega_3(D) + \#\Omega_2^-(D) + \#\Omega_T^-(D) + \#\Omega_1^-(D). \tag{3-1}$$

Note that multiedges do not create issues in the sum, as they are counted separately.

It follows from (3-1) that the valence of any diagram is bounded from above by quantities depending only on the knot projection:

$$v(D) \leq 8\alpha(D) + p_1(D) + p_2(D) + p_3(D) + 2\sum_{k\geq 2} p_k(D)k(k-1). \tag{3-2}$$

Equation (3-2) is obtained by giving an upper bound on the possible Ω^- and Ω_3 moves in terms of the number of edges of the regions interested by the moves (ie on the number of 1–, 2– and 3–regions for Ω_1^-, $\Omega_2^- - \Omega_T^-$ and Ω_3 moves respectively).

[10] In the present discussion we find it convenient to blur a bit the distinction between Ω_2 and Ω_T, since we are only interested in the total count.

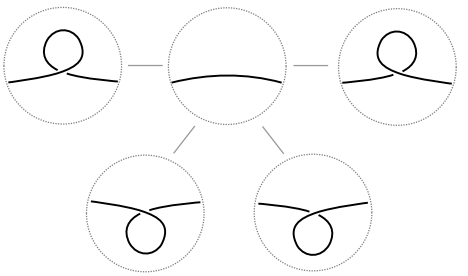

Figure 35: The possible Ω_1^+ moves that can be performed on each arc.

Looking at (3-1) we can obtain a lower bound as well, which allows to say that the valence grows at least linearly with the crossing number. Define $P(D)$, the maximal period of a nontrivial diagram D, as the maximal order of a finite group acting on the sphere (or the plane), preserving the diagram setwise.[11] Recall that if K is not the unknot, then K admits finitely many orders of periodicity (see [7, Theorem 3]).

Lemma 3.10 *If D is a nontrivial knot diagram with periodicity $P(D)$* (where $P(D)$ is equal to 1 if D is nonperiodic), *then*

$$v(D) \geq \frac{8\alpha(D)}{P(D)}. \tag{3-3}$$

This follows easily by observing that each fundamental domain for the periodic action must contain at least one arc.

Of course if D is nonperiodic, the lower bound

$$v(D) \geq 8\alpha(D) + \sum_{k\geq 2} p_k(D)k(k-1) + k_{\text{ext}}(k_{\text{ext}}-1) \tag{3-4}$$

holds as well.

Proposition 3.11 *The minimal valence δ detects the unknot $\bigcirc$.*

Proof $\delta(\bigcirc) = 3$, as shown in Figure 37, while if $K \neq \bigcirc$, then for every diagram D representing K we have $v(D) \geq 4$, since each fundamental domain for a periodic action must contain at least one arc (as in the proof of Lemma 3.10), and for every arc there are at least 4 (two Ω_T^+ and two Ω_1^+) possible moves. □

Lemma 3.12 *For each knot K, the number of vertices in $\mathcal{G}_P(K)$ or $\mathcal{G}(K)$ whose valence is bounded by a constant is finite.*

[11] We need to exclude the trivial diagram of the unknot to ensure that $P(D)$ is in fact finite.

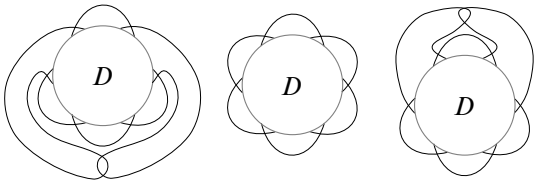

Figure 36: The two nonequivalent possibilities for an Ω_2 move in the external zone.

Proof This follows immediately from the fact that there are only finitely many diagrams of a knot with crossing number bounded by a constant, finitely many periods for each knot, and by (3-3) the valence is bounded from below by a linear function in $\mathrm{cr}(D)$. □

In particular, choosing $\delta(K)$ as the constant in the previous lemma, we get:

Corollary 3.13 *$\#\delta(K)$ is well defined.*

Following [12], we call a diagram $D \in \mathcal{D}(K)$ *hard* if

$$\#\Omega_1^-(D) = \#\Omega_2^-(D) = \#\Omega_T^-(D) = \#\Omega_3(D) = 0.$$

We can refine (3-1) for hard diagrams:

Corollary 3.14 *If D is a hard diagram of a nonperiodic knot K, then*

$$v_S(D) = 8\alpha(D) + \sum_{k>1} k(k-1)p_k(D).$$

The analogous result for $\mathcal{G}_P(K)$ is obtained by adding $k_{\mathrm{ext}}(k_{\mathrm{ext}}-1)$.

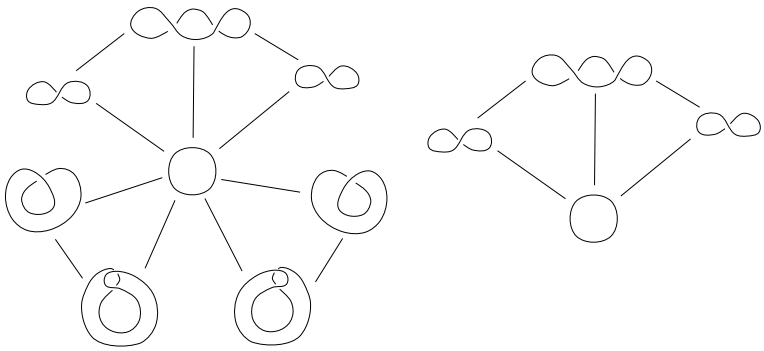

Figure 37: The ball $S(\bigcirc)$ in the planar (left) and S^2 (right) graphs.

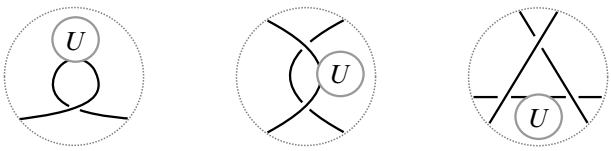

Figure 38: How to kill Reidemeister moves.

In [12] Kauffman and Lambropoulou exhibit an infinite family of hard unknots. Using their result, it is not difficult to argue that every knot admits (infinitely many) hard diagrams. Take any diagram $D \in \mathcal{D}(K)$, and choose a (nontrivial) hard diagram U of the unknot. If D is not hard, choose an Ω_i^- or Ω_3 move and perform a diagram connected sum with U to "kill" it as in Figure 38. Generally, hard diagrams of nonperiodic knots are interesting from the $\mathcal{R}$–graphs viewpoint, since for them the valence is completely determined by the knot projection, rather than by the diagram. In particular this implies that given a hard diagram, it will have minimal valence among all the diagrams obtained from it by changing any number of crossings.[12]

Remark 3.15 It is possible to compute the valence of the two trefoil knots of Figure 2 in $\mathcal{G}_P(3_1)$. Taking into account the periodicities of the two diagrams (it is of order 3 for the first and 2 for the other), one gets that (as planar diagrams) the first has valence 24 and the second 32, so they are set apart in $\mathcal{G}_P(3_1)$. The valence in $\mathcal{G}(3_1)$ instead is 12. We will in fact prove in an upcoming paper that $\delta_P(3_1) = 24$ and $\delta(3_1) = 12$, and that in both cases $\#\delta(3_1) = 1$.

In order to facilitate the proof of Theorem 1.1, understanding how the valence of a diagram can change under the various Reidemeister moves is crucial.

It is of course impossible to a priori compute the difference of the valence between two vertices at distance 1, since this value depends on the crossings and specific configurations in the diagrams involved. It is, however, possible to pinpoint a quite good bound by accounting for the number of edges of the regions interested by the Reidemeister move.

This last task is a quite tedious exercise; in the following we denote by[13] $\varepsilon_{j,i}$ and $\varepsilon_{j,3}$ the difference in the number of Reidemeister moves of type Ω_i^- and Ω_3, respectively,

[12] This is no longer true if one of the diagrams obtained by changing some crossings in a hard one is periodic.

[13] We suppress the dependency of the $\varepsilon_{j,i}$ from the diagrams in the notation for aesthetic reasons.

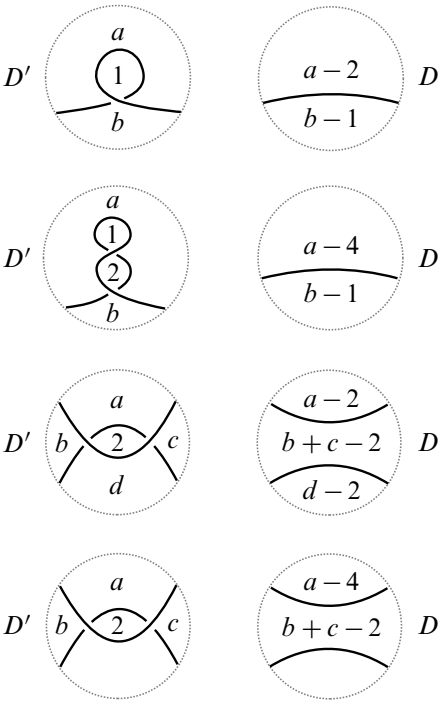

Figure 39: The four possible configurations considered in Propositions 3.16–3.19.

that can be performed on two diagrams differing by a single Reidemeister move Ω_j^+ with $i, j \in \{1, T, 2\}$.

If $D' = \Omega_1^+(D)$, then

- $\varepsilon_{1,1} = \#\Omega_1^-(D') - \#\Omega_1^-(D) \in \{0, 1\}$;
- $\varepsilon_{1,T} + \varepsilon_{1,2} = (\#\Omega_T^-(D') - \#\Omega_T^-(D)) + (\#\Omega_2^-(D') - \#\Omega_2^-(D)) \in \{-2, 0, 1\}$;[14]
- $\varepsilon_{1,3} = \#\Omega_3^-(D') - \#\Omega_3^-(D) \in \{-4, \dots, 4\}$.

We denote the sum of the ε contributions in each case as $\sum_i \varepsilon_{1,i}$; these count the part of the valence of a diagram that is not completely determined by the knot projection. In particular, we have that

$$-6 \leq \sum_{i \in \{1,2,T,3\}} \varepsilon_{1,i} \leq 6. \tag{3-5}$$

[14] Here we consider the sum $\varepsilon_{1,T} + \varepsilon_{1,2}$ since performing an Ω_1^+ move at the top of a preexisting tentacle may decrease the number of Ω_T moves, changing them in Ω_2 moves.

Proposition 3.16 *If K is a nonperiodic knot and $D' \in \mathcal{D}(K)$ is obtained from D by adding a curl* (ie performing an Ω_1^+ move, as in the upper part of Figure 39), *then $v(D') > v(D)$. More precisely, if the move is* ***internal****, that is, the two zones involved are not the external one, then*

$$v(D') = v(D) + 8 + 4a + 2b + \sum_j \varepsilon_{1,j}. \tag{3-6}$$

If the zone with a edges is external, then

$$v(D') = v(D) + 2 + 8a + 2b + \sum_j \varepsilon_{1,j}. \tag{3-7}$$

And finally if the zone with b edges is external, then

$$v(D') = v(D) + 6 + 4a + 4b + \sum_j \varepsilon_{1,j}. \tag{3-8}$$

Moreover, we have $\sum_j \varepsilon_{1,j} \in \{-6, \dots, 6\}$. Thus, performing an Ω_1^+ move always increases the valence.

Proof After performing an Ω_1^+ move, the number of arcs in D' increases by 2, that is, $\alpha(D') = \alpha(D) + 2$. Moreover, assuming that a, b, $a-2$ and $b-1$ are pairwise distinct, we have the following changes in the p_k:

$$\begin{aligned} p_a(D') &= p_a(D) + 1, \\ p_{a-2}(D') &= p_{a-2}(D) - 1, \\ p_b(D') &= p_b(D) + 1, \\ p_{b-1}(D') &= p_{b-1}(D) - 1. \end{aligned}$$

Adding everything up, and keeping in mind (3-1), we obtain

$$v(D') - v(D) = 8 \cdot 2 + a(a-1) - (a-2)(a-3) + + b(b-1) - (b-1)(b-2) + \sum_j \varepsilon_{1,j}.$$

That is, precisely

$$v(D') = v(D) + 8 + 4a + 2b + \sum_j \varepsilon_{1,j}.$$

Notice that even if a, b, $a-2$ and $b-1$ are not pairwise distinct, the same computation holds. All other Ω_j^- moves (that do not depend solely on the knot projection) add up to $\sum_j \varepsilon_{1,j}$.

To obtain (3-7) and (3-8) it is enough to add the contribution of the external region, which is $a(a-1) - (a-2)(a-3) = 4a - 6$ in the first case, and $b(b-1) - (b-1)(b-2) = 2b - 2$ in the second. □

The proof is identical in the other cases considered below, and we are going to omit it.

Proposition 3.17 *Let $D, D' \in \mathcal{D}(K)$ be two nonperiodic diagrams differing by an Ω_T creating a tentacle of length 1* (as in the upper-middle part of Figure 39)*. Then, if the zones involved are not external,*

$$v(D') = v(D) + 12 + 8a + 2b + \sum_j \varepsilon_{T,j}. \tag{3-9}$$

If the zone with a edges is external,

$$v(D') = v(D) - 8 + 16a + 2b + \sum_j \varepsilon_{T,j}. \tag{3-10}$$

And finally if the zone with b edges is external,

$$v(D') = v(D) + 10 + 8a + 4b + \sum_j \varepsilon_{T,j}. \tag{3-11}$$

Proposition 3.18 *If two nonperiodic diagrams $D, D' \in \mathcal{D}(K)$ differ by an Ω_2 move in which the regions with a and d edges do not coincide* (as in the middle part of Figure 39)*, then if the move is internal,*

$$v(D') = v(D) + 16 + 4(a + b + c + d) - 2bc + \sum_j \varepsilon_{2,j}. \tag{3-12}$$

Proposition 3.19 *If two nonperiodic diagrams $D, D' \in \mathcal{D}(K)$ differ by an Ω_2 move in which the regions with a and d edges coincide* (as in the lower part of Figure 39)*, then if the move is internal,*

$$v(D') = v(D) + 8 + 4(2a + b + c) - 2bc + \sum_j \varepsilon_{2,j}. \tag{3-13}$$

Remark 3.20 An Ω_T creating a tentacle of length greater than 1 is a special case of Proposition 3.19 in which $c = 2$. Thus, in this case we obtain

$$v(D') = v(D) + 8 + 4(2a + b + 2) - 4b + \sum_j \varepsilon_i = v(D) + 16 + 8a + \sum_j \varepsilon_{T,j}.$$

It is worth remarking that, when dealing with $\mathcal{G}(K)$, the change of the valence is determined by (3-6), (3-9) and (3-12) in the respective cases.

Remark 3.21 We will find it useful to divide the valence of every vertex in two parts, namely the positive valence $v^+(D)$ and the negative valence $v^-(D)$. The positive valence is defined as the number of edges emanating from D which correspond

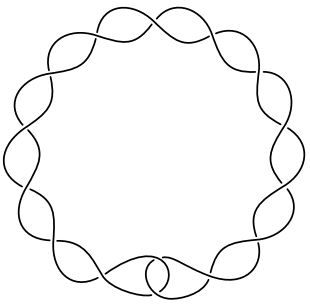

Figure 40: By performing an Ω_2^+ move in the central (or external) region, we obtain a diagram with lower valence.

to Ω_*^+ moves, where $* \in \{1, 2, T\}$. Note that $v^+(D)$ only depends on the projection of D. If we wish to consider only the positive valence, (3-6), (3-9), (3-12) and (3-13) can be rewritten as

$$v^+(D') = v^+(D) + 8 + 4a + 2b, \tag{3-14}$$

$$v^+(D') = v^+(D) + 12 + 8a + 2b, \tag{3-15}$$

$$v^+(D') = v^+(D) + 16 + 4(a + b + c + d) - 2bc, \tag{3-16}$$

$$v^+(D') = v^+(D) + 8 + 4(2a + b + c) - 2bc. \tag{3-17}$$

Remark 3.22 Proposition 3.17 suggests how to produce examples of knots in which the minimal complexity is not realised by a diagram minimising the crossing number. From (3-12) it is apparent that if b or c is sufficiently large, then the diagram D' (with higher crossing number than D) obtained by performing an Ω_2^+ move will have a lower valence. An easy example of this phenomenon is given in Figure 40. This is the twist knot with 17 crossings; note that the example shown is also alternating, reduced and nonperiodic. When the internal (and external) region has more than 12 faces, performing an Ω_2^+ move decreases the valence, according to (3-12) (with $a = d = 4$ and $b = c = 8$).

In particular, any knot in which all diagrams realising the crossing number have many regions with a sufficiently high number of edges provides an example where the minimal valence is not realised in the diagram with minimal crossing number.

We prove here some facts that are going to be useful in the next sections. First of all we show that the S^2–graph can *distinguish* between the different Reidemeister moves. This means that by looking at a neighbourhood of an edge of $\mathcal{G}(K)$, we can

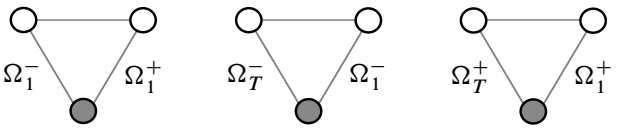

Figure 41: The grey dots represent the diagram D we start from.

tell which Reidemeister move it represents; furthermore this will provide a way to read the crossing number of a diagram D from the combinatorial structure of $S(D)$.

Theorem 3.23 *The S^2–graph distinguishes the Reidemeister moves and detects the crossing number of a diagram.*

Proof In the interest of clarity we are going to start by examining the nonperiodic case. Fix a diagram $D \in \mathcal{D}(K)$ for a nonperiodic knot K. The combinatorics of $S(D)$ will allow us to distinguish the various moves.

Since by Theorem 3.2 all Ω_1 moves are paired with at least one Ω_T move in a triangle, it is easy to argue that the graph can tell apart the two sets of moves $M_1 = \{\Omega_1^{\pm}, \Omega_T^{\pm}\}$ and $M_2 = \{\Omega_2^{\pm}, \Omega_3\}$.

To further separate the elements of M_1 we can thus restrict to triangles in $S(D)$. Choose an edge emanating from a vertex D which is part of a triangle. There are three possibilities, shown in Figure 41.

From this, using Proposition 3.16, it is easy to argue that $\mathcal{G}(K)$ can tell apart the elements in M_1; indeed, if only one of the two moves decreases the valence, then they are both Ω_1 moves, and the one which decreases it is the Ω_1^-. If both moves decrease the valence, then the one that decreases it most is the Ω_T^-, and the other is an Ω_1^-. Lastly, if both moves increase the valence, then the one that increases it most is an Ω_T^+, and the other is an Ω_1^+.

Now, since the number of Ω_1^+ moves is a multiple of the arc number of the diagram (see Figure 35), the crossing number $\mathrm{cr}(D)$ corresponds to $\frac{1}{8}(\#\Omega_1^+(D))$. Hence, since we can distinguish and count such moves, we can read the crossing number of a diagram from $S(D)$.

Using this information we can tell apart the elements of M_2 as well and conclude that the only remaining moves are $\Omega_2^{\pm}$ moves and Ω_3 moves, all of which are not part of a triangle. These appear as edges connected only to the centre of $S(D)$. We

can distinguish between them by counting the crossing number of the vertices they connect D to; one then just needs to recall that Ω_3 moves do not increase it, while Ω_2 increase or decrease it by 2. Hence it follows that we can distinguish among Ω_2^+, Ω_2^- and Ω_3 moves as well.

In all the previous discussion, in order to determine $\mathrm{cr}(D)$, we only used the fact that all diagrams in $S(D)$ were nonperiodic; this fact will allow us to compute it in the periodic case as well.

If K is periodic we cannot use directly the various equations relating the valence of two neighbouring vertices, since one[15] could be periodic.

Instead of trying to directly detect from the structure of the graph whether a diagram is periodic, we can use Lemma 3.7 to bypass most complications. For every vertex D, define the generalised triangle number by $n_{\mathrm{tr}}(D) = \#\Omega_1^{\pm}(D) + \#\Omega_T^{\pm}(D)$. This quantity is computable from the graph, since by Theorem 3.2 it coincides with the number of edges emanating from D which are part of at least one triangle. By Lemma 3.7, at least two diagrams appearing in a triangle, reached by an Ω_1^+ and an Ω_T^+ respectively, will be nonperiodic, and we claim that such diagrams maximise n_{tr} among all the diagrams reached by edges starting from D that are part of at least one triangle. Let $a_h(D)$ denote the number of (maximal) height h tentacles in D, and define

$$n(D) = p_1(D) + \sum_{h \geq 1} h a_h(D).$$

Note that $n(D)$ is equal to the sum $\Omega_1^-(D) + \Omega_T^-(D)$ when D is not periodic. Then, for the diagrams D'' and D' in Lemma 3.7,

$$\begin{aligned} n_{\mathrm{tr}}(D'') &= 16\,\mathrm{cr}(D'') + n(D'') = 16(\mathrm{cr}(D) + 2) + n(D) + 2, \\ n_{\mathrm{tr}}(D_1') &= 16\,\mathrm{cr}(D') + n(D') \ = 16(\mathrm{cr}(D) + 1) + n(D) + 1. \end{aligned}$$

This follows since we are performing the curls on the top of a tentacle (or on any arc, if there are no curls in D), and this fact ensures that the number of Ω_*^- moves is equal to $n(D) + 1$ when $* = 1$ and to $n(D) + 2$ when $* = T$. On the other hand, for any diagrams D_T^+ and D_1^+ reached from D by an Ω_T^+ move and an Ω_1^+ respectively,

$$\begin{aligned} n_{\mathrm{tr}}(D_T^+) &\leq 16(\mathrm{cr}(D) + 2) + n(D) + 2, \\ n_{\mathrm{tr}}(D_1^+) &\leq 16(\mathrm{cr}(D) + 1) + n(D) + 1. \end{aligned}$$

[15] Or both; see also Question 6.5.

The presence of periodicity in D_1^+ or D_T^+ can only decrease the value of n_{tr}, and the same holds if the moves are not performed (with the appropriate sign) on the top of a preexisting tentacle. In other words these moves maximise $\#\Omega_1^- + \#\Omega_T^-$. If we consider moves that decrease the crossing number, disregarding the possible periodicities, the numbers n_{tr} we obtain have no chance of being greater than $n_{\mathrm{tr}}(D'')$. So, choose the diagrams in $S(D)$ maximising this quantity; they correspond to vertices reached by Ω_T^+ moves. Consider all the edges that form triangles with them: these have to correspond to diagrams reached by Ω_1^+ moves. Choose between them one maximising n_{tr}. Notice that D is nonperiodic if and only if $n_{\mathrm{tr}}(D) = n_{\mathrm{tr}}(D') - 17$. Now, choose D''' in $S(D')$ forming a triangle with D'', with $S(D''')$ totally nonperiodic, and such that it maximises n_{tr} in $S(D')$. We know that such a diagram exists by Lemma 3.7, and we can check the hypothesis on the nonperiodicity of $S(D''')$ thanks to the above criteria. Then, we can recover $\mathrm{cr}(D''')$, and obtain $\mathrm{cr}(D)$ as $\mathrm{cr}(D''') - 3$.

Hence, using the crossing number as in the nonperiodic case, we can tell apart the various types of moves, and we are done. □

This last result will allow us to say "perform an $\Omega_i^{\pm}$ move on a diagram" in a way that is meaningful also at the level of the graph. In other words, we just proved that the $\mathcal{R}$–graphs intrinsically contain the same amount of information as the same graphs with edges decorated according to which $\Omega_i^{\pm}$ move we are performing.

By the previous result we know that the crossing number of a diagram can be read by looking at $S(D)$. Thus if a knot is nonperiodic, taking the minimum of $\frac{1}{8}(\#\Omega_1^+)$ among all vertices of the corresponding Reidemeister graphs gives back $\mathrm{cr}(K)$, the crossing number of the knot. For periodic knots, this procedure produces a slightly different invariant, which can be regarded as crossing number up to periodicities. More precisely, define

$$\widehat{\mathrm{cr}}(K) = \tfrac{1}{8} \min_{D \in \mathcal{D}(K)} \#\Omega_1^+(D).$$

If K is nonperiodic, $\widehat{\mathrm{cr}}(K) = \mathrm{cr}(K)$, while in general $\widehat{\mathrm{cr}}(K) \leq \mathrm{cr}(D)$. As an example, we have $\widehat{\mathrm{cr}}(3_1) = \frac{1}{2}$.

Note that a similar consideration for the other kinds of moves does not yield useful invariants: it is possible to show that the minimal number of Ω_2^+ moves is simply related to the combinatorics of the number of regions in the complement of the diagram on $\mathbb{R}^2$ or S^2, and the minimal number of Ω_2^- and Ω_3 moves one can perform within a knot type is always 0 (as can be seen by "killing" all the Ω_3 moves with an Ω_1 in

the region with three edges, similarly to what was done in Figure 38). Nonetheless one might obtain some meaningful invariants by restricting diagrams not minimising the valence.

The knowledge of the crossing number from the graph also implies that we can use a result of Coward and M Lackenby [5] to give some upper bounds on the path distance between two diagrams.

4 Global properties

This section is devoted to the analysis of some global properties of the $\mathcal{R}$–graphs. We begin by proving that each $\mathcal{R}$–graph is not hyperbolic.

Proposition 4.1 *The $\mathcal{R}$–graphs are not hyperbolic.*

Proof Choose a nonperiodic diagram $D \in \mathcal{D}(K)$ containing no 1–regions, an arc on D, and a polygon P having this arc as a face. We can embed isometrically the rank 2 lattice graph as follows: to the pair $(a, b) \in \mathbb{Z}^2$ associate the configuration on the arc composed of a positive curls in the region P if $a > 0$, and in the other region touching the arc if $a < 0$; do the same for b, this time with negative crossings on the right of the previous ones. An example is shown in Figure 42. The fact that the embeddings are isometric follows, for example, from an analysis of the I_{lk} invariants of the diagrams: $I_{\mathrm{lk}}(D_{a,b}) = I_{\mathrm{lk}}(D) + aX_0 + bY_0$, where $D_{a,b}$ is the diagram corresponding to the element (a, b). □

Proposition 4.2 *The Reidemeister graphs $\mathcal{G}_P(K)$ and $\mathcal{G}(K)$ are not planar.*

Proof We are going to prove that for every knot K we can find a K_5 minor[16] contained in each $\mathcal{R}$–graph of K. This is achieved by considering the local construction shown in Figure 43. The edges denoted with a Greek letter are length 2 paths; as shown in Figure 44, these can be obtained by putting the two moves alongside each other and then resolving either one. □

In graph theory, it is customary to consider *coarse* properties of a (infinite and locally finite) graph. One way to do this is to study the quasi-isometry class of the graph, often through related invariants.

[16] As is customary, K_n denotes the complete graph on n vertices.

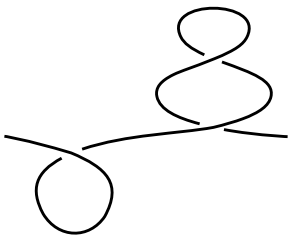

Figure 42: This configuration represents $(-1,2) \in \mathbb{Z}^2$.

A *ray* of a locally finite graph G is a semi-infinite simple path in G; two rays $r_1, r_2 \subset G$ are regarded as equivalent if there exists a third ray r_3 containing infinitely many vertices of both r_1 and r_2.

An *end* is an equivalence class of rays, and it is called *thick* if it contains infinitely many pairwise disjoint rays.

Proposition 4.3 *Each S^2–Reidemeister graph has only one thick end.*

Proof It is immediate to show (eg using paths such as those in Figure 45 or tentacle configurations) that there are infinitely many disjoint rays in $\mathcal{G}_P(K)$ and $\mathcal{G}(K)$ for each choice of $K \in \mathcal{K}$. To show that there is only one end, we will prove that removing any ball with arbitrary radius does not disconnect the graphs into two pieces, each containing infinitely many vertices. This in turn would immediately imply that there is only one equivalence class of rays in the graph.

Consider a diagram $D \in \mathcal{D}(K)$ for a knot K, and the radius R ball $S_R(D)$ in $\mathcal{G}(K)$ (or equivalently in $\mathcal{G}_P(K)$). Call H the maximal height among the tentacles of the

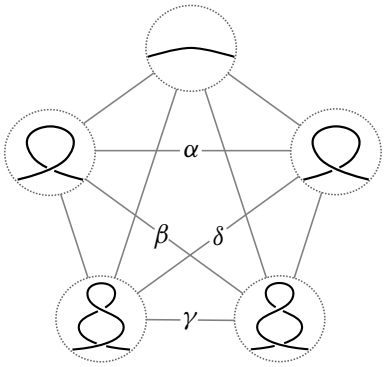

Figure 43: A local embedding of K_5 as a minor of any $\mathcal{G}(K)$.

Figure 44: The path corresponding to the β edge in Figure 43.

diagrams contained in $S_R(D)$, and take any two diagrams $D_0, D_1 \in \mathcal{D}(K)$ which do not belong to $S_{R+H+1}(D)$; we need to find a path in $\mathcal{G}(K) \setminus S_R(D)$ connecting D_0 to D_1. Choose an arc on D_0 and on D_1, and create on each a tentacle of height greater than H. These two new diagrams D_0' and D_1' can be connected through moves that avoid the newly created tentacles,[17] and this path γ from D_0' to D_1' will not intersect $S_R(D)$, thanks to the hypothesis on H. Attaching to the ends of γ the two paths $\hat{\gamma}_i$ from D_i to D_i' induced by the creation of the tentacles gives the desired path from D_0 to D_1. Note that the hypothesis on the height of the tentacle allows us to say that the paths $\hat{\gamma}_i$ do not intersect $S_R(D)$, and $S_{R+H+1}(D) \setminus S_R(D)$ contains only finitely many vertices. □

The S^2–Reidemeister graphs contain only one end, but infinitely many disjoint rays, hence by Halin's grid theorem [9], each must contain a subdivision of the planar hexagonal tiling.

Remark 4.4 One might find it reasonable to assume that the graphs $\mathcal{G}_P(K)$ and $\mathcal{G}(K)$ are quasi-isometric; it is, however, easy to see that the "natural" map[18] $\mathcal{G}_P(K) \hookrightarrow \mathcal{G}(K)$ between the two graphs fails to be a quasi-isometry. This can be seen from Figure 36: the two diagrams on the left and right can have arbitrarily large distance in $\mathcal{G}_P(K)$, but are identified in $\mathcal{G}(K)$.

These graphs also exhibit a fractal behaviour, which can be observed, for example, by considering sequences of Ω_1 moves as in Figure 45. The corresponding subgraph can be embedded (infinitely many times) in each $\mathcal{R}$–graph for any knot K.

The $\mathcal{R}$–graphs can be filtered in several ways; the easiest one is to consider the filtration induced by the distance from the vertices with minimal valence.

Given a knot K, denote by $\mathcal{F}_m(K)$ the subgraph spanned by those vertices whose distance from the minimal diagrams of K is $\leq m$, and denote by $\#\mathcal{F}_m(K)$ the number of vertices it contains.

[17] Remember that we are working on S^2.

[18] Here we map a planar diagram D to its equivalence class in $\mathcal{G}(K)$.

We can extract some numerical invariants from this filtration on $\mathcal{G}(K)$:

Definition 4.5 Define

$$f_K\colon \mathbb{N} \to \mathbb{N}, \quad f_K(m) = \#\mathcal{F}_m(K),$$

and

$$M(K) = \min_{m \geq \delta(K)} \{\pi_0(\mathcal{F}_m(K)) = \mathbb{Z}\}.$$

In other words, $M(K)$ measures the minimal distance between the diagrams of minimal complexity in $\mathcal{G}(K)$. In particular, $M(K) = 0$ if and only if a knot type is simple.

Recalling the proof of Lemma 3.12, we can also define another filtration on $\mathcal{G}(K)$:

Definition 4.6 Let $\widetilde{\mathcal{F}}_K$ be the filtration of $\mathcal{G}(K)$ whose m–level consists of the vertices of $\mathcal{G}(K)$ with valence less or equal to m. Let also

$$g_K\colon \mathbb{N} \to \mathbb{N}$$

be defined as the associated counting function

$$g_K(m) = \#\{D \in \mathcal{G}(K) \mid v(D) \leq m\}.$$

Clearly $g_K(m) = 0$ for all $m < \delta(K)$, and $g_K(\delta(K)) = \#\delta(K)$.

Both these filtrations $\mathcal{F}$ and $\widetilde{\mathcal{F}}$, together with the associated integer-valued counting functions f_K and g_K, are knot invariants, and it is not hard to show that they both distinguish the unknot. Moreover one can consider the homology groups of the various level sets and obtain yet other knot invariants.

In [15], Miyazawa computes the homology groups of the *Reidemeister complex*, which he denotes by $M(K : P_5, 1)$, in the case of oriented diagrams with a minimal generating set of Reidemeister moves. Along these lines we can define a slightly different version of Reidemeister complex, denoted by $\mathcal{CG}(K)$, as follows: an n–simplex $\Delta_n = \langle D_0, \dots, D_n\rangle$ is given by a string of $n+1$ distinct diagrams such that[19] $d(D_i, D_j) = 1 - \delta_{i,j}$, considered up to permutations of the indices.

Define $C_n(\mathcal{CG}(K))$ as the free abelian group generated by n–simplices, with the obvious boundary operator induced by simplicial homology,

$$\partial(\langle D_0, \dots, D_n\rangle) = \sum_{i=0}^{n} (-1)^i \langle D_0, \dots, \widehat{D_j}, \dots, D_n\rangle. \tag{4-1}$$

[19] Here $\delta_{i,j}$ denotes Kronecker's delta function, and d is the path distance.

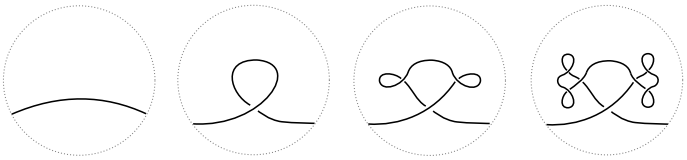

Figure 45: The "fractal behaviour" of $\mathcal{G}(K)$.

From this perspective, $\mathcal{G}(K)$ is the 1–skeleton of $\mathcal{CG}(K)$. Miyazawa proved that $H_0(M(K:P_5,1);\mathbb{Z}) = \mathbb{Z}$ (which follows from Reidemeister's theorem), and that $H_n(M(K:P_5,1);\mathbb{Z}) = 0$ for every $n \geq 2$ and $K \in \mathcal{K}$.

Our situation is slightly different; with the methods developed in Section 2 we can easily establish the triviality of $H_n(\mathcal{CG}(K);\mathbb{Z})$ for $n \geq 3$:

Proposition 4.7 *For any knot K we have $H_n(\mathcal{CG}(K);\mathbb{Z}) = 0$ for $n \geq 3$* (in both the planar and S^2 cases).

Proof Assume there is a tetrahedron Δ_3 in $\mathcal{CG}(K)$; then it follows from Theorem 3.2 that all faces have to be composed of triples $\Omega_T^{\pm} - \Omega_1^{\mp} - \Omega_1^{\mp}$. Up to symmetries, there is only one possibility to be considered, shown in Figure 46. However this can be excluded as well, by taking into account the signs of the moves composing the tetrahedron. In particular this shows that there are no simplices of dimension $n \geq 3$, hence all the corresponding homology groups vanish. □

In particular it follows that $\mathcal{CG}(K)$ is just $\mathcal{G}(K)$ with all triangles capped by 2–simplices. It is not hard to prove that $H_1(\mathcal{CG}(K);\mathbb{Z})$ is an infinitely generated free abelian group, as any pair of distant $\Omega_3 - \Omega_2$ moves does not bound any union of 2–simplices. We can now conclude the computation of the homology groups of $\mathcal{CG}(K)$:

Proposition 4.8 $$H_2(\mathcal{CG}(K);\mathbb{Z}) \cong \mathbb{Z}^{\infty}.$$

Proof The two configurations in the top part of Figure 27 show a topologically embedded 2–sphere in $\mathcal{CG}(K)$. As these are local configurations, they can be found infinitely many times on any $\mathcal{G}(K)$. □

So, as in Miyazawa's case, the global homology does not provide useful invariants. Nonetheless, some properties of the diagrams can be inferred from the local homology of the complexes. Denote by $S^{\mathrm{cpx}}(D)$ the ball of radius 1 centred in D, seen as a subcomplex of $\mathcal{CG}(K)$.

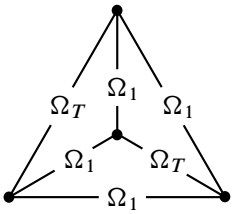

Figure 46: The only tetrahedron with compatible faces. There is no way of coherently orienting the signs on the edges of its faces.

Lemma 4.9 *If $H_1(S^{\mathrm{cpx}}(D);\mathbb{Z}) = \mathbb{Z}^m$, then $m \geq \widehat{\mathrm{cr}}(D)$. Also, $H_2(S^{\mathrm{cpx}}(D);\mathbb{Z}) = 0$ if and only if $p_1(D) = 0$.*

The first part of this lemma follows easily from Figure 8, while the second is a consequence of Lemma 3.5 together with Theorem 3.2.

5 Completeness of the S^2–graph invariant

This section is devoted to the proof of Theorem 1.1 (recalled below). The proof will rely solely on results from Section 3, and exploits rather large portions of the graph.

Theorem 1.1 *The S^2–Reidemeister graph is a complete knot invariant up to mirroring; that is, $\mathcal{G}(K) \equiv \mathcal{G}(K')$ if and only if K' is isotopic to K or $\overline{K}$.*

We will prove Theorem 1.1 by breaking it down in smaller parts, which are the content of the following propositions.

Suppose we have a knot $K \in \mathcal{K}$, and suppose that $D \in \mathcal{D}(K)$ is any diagram. Write $P(D) = (p_1(D), \dots, p_m(D))$, where m is the greatest coefficient with a nonzero entry (or equivalently the maximal number of sides among the regions in the complement of D in S^2).

Proposition 5.1 *The S^2–graph of a knot determines $P(D)$ for each vertex D such that all diagrams in $S(D)$ are nonperiodic and $p_1(D) = 0$.*

Proof Thanks to Lemma 3.5 we know that if a diagram D does not contain any curls, then all the triangles in $S(D)$ which admit D as the vertex with lower crossing number are normal. Moreover, since there are $8\,\mathrm{cr}(D)$ Ω_1^+ moves and Ω_T^+ moves, we can conclude that in $S(D)$ there are exactly $8\,\mathrm{cr}(D)$ triangles. For each Ω_1^+ move, choose

the corresponding Ω_T^+ move; call D' and D'' respectively the diagrams obtained by performing these Ω_1^+ and Ω_T^+ moves on D. By (3-14) and (3-15), the difference of the positive valences is

$$v^+(D'') - v^+(D') = 4a + 20,$$

where a is the number of edges of the region in which the tentacle and the curl will appear.

If we do the same for all possible Ω_1^+ moves applicable to D, we get a multiset of numbers $\{n_i'\}_{i\in\{0,\dots,\#\Omega_1^+(D)\}}$, where each entry is the difference in v^+ for a given Ω_1^+; define a new multiset $\{n_i\}_{i\in\{0,\dots,\#\Omega_1^+(D)\}}$, where $n_i = \frac{1}{4}(n_i' - 20)$. Each region with a sides contributes to this new list with exactly[20] $2a$ entries equal to a. It is thus immediate to show that we can compute each $p_a(D)$ from $S(D)$. □

However the knowledge of $P(D)$ on a subset of vertices does not immediately guarantee the completeness of $\mathcal{G}(K)$. A priori there might be two distinct knots (up to mirroring) whose diagrams have the same number and types of Reidemeister moves and such that their complement has the same number of regions. We first need to detect the structure of D as a 4–valent graph on S^2.

Proposition 5.2 *The S^2–graph recognises the projections of the knot corresponding to diagrams without 1–regions.*

Proof Let us deal only with nonperiodic knots for the moment, and come back to the periodic case afterwards.

Choose a vertex $D \in \mathcal{D}(K)$ with $p_1(D) = 0$. This is possible by Theorem 3.23, since the condition $p_1(D) = 0$ is equivalent to the absence of Ω_1^- moves emanating from D.

To obtain the structure of D as a graph, we need to be able to tell which regions are adjacent to one another in S^2, or in other words, we want to determine the dual graph of the projection.

We need to look for this information outside of $S(D)$; begin by assuming for the moment that D only has one region R with a certain number k of edges (that is, $p_k(D) = 1$), and k is such that there are no regions with $k \pm l$ sides, for any $l \leq L$ where L is a suitably big integer.

[20] Corresponding to the two possible Ω_1^+ moves performed with the curl contained in the region on one of its edges.

To determine the number of edges of the regions adjacent to R, perform an Ω_1^+ move on one edge[21] of R so that the curl is contained in the interior of R. We can then compute the number of edges of the other region involved in the move as follows. The Ω_1^+ move is associated to a unique Ω_T^+ move, connecting D with a diagram D'' with which they form a triangle. By counting the difference of the positive valences between D' and D'' and by using (3-14) and (3-15) we can compute the number of edges of R. Once we have that, the difference in the positive valence between D and D' gives us the number of edges of the other region involved. Moreover, note that knowing the number of edges in these two regions is enough to compute $P(D')$ from $P(D)$.

Repeating for all edges[22] in R we get the number of sides of each region which shares an edge with R.

From this last paragraph it follows that if we could find a diagram of K such that, with the exception of regions with one side, all the regions have a different and sufficiently spaced number of sides, then we could infer how they are globally "patched together" to form the corresponding 4–valent graph.

There is an easy way to achieve such a configuration in a controlled way. Start with a diagram D with $p_1(D) = 0$, and perform an Ω_1^+ move; again by the previous line of thought[23] we can tell that the move has been made with the curl contained in a region with a edges which is adjacent to a region with b edges (as in the top of Figure 39).

Call D_1 the diagram obtained; we can recover $P(D_1)$ from $P(D)$, since we know the number of edges of the regions involved. There are only two possible choices to perform another Ω_1^+ move on this new diagram in such a way that an Ω_1–multiedge with valence 2 is created (given by performing an identical Ω_1^+ move on the left or right of the previous one). We can repeat this process $N_1 \gg 0$ times, obtaining a new diagram D_{N_1} with only one region with $a + 2N_1$ edges and a region with $b + N_1$ edges and such that there is only one multiedge of order N_1 and exactly N_1 regions with one edge. Notice again that at each step we can recover $P(D_i)$ from $P(D_{i-1})$,

[21] This is in fact well defined on the graph, since the valence of the diagrams obtained in this fashion will be different from any other obtainable by making an Ω_1^+ anywhere else.

[22] Thanks to the hypothesis on R, we can recognise from the graph all the Ω_1^+ which create a curl in R.

[23] Since triangles in $S(D)$ are normal, we can compute $P(D)$ by considering the difference in the valence between diagrams reached by triangles. Moreover, by considering one triangle at the time, and using the differences in the positive valences between D'' and D', and between D' and D, we can compute the number of edges involved in the corresponding Ω_1^+.

since at each step we already know the number of edges of one region involved, and from the difference in the positive valence between D_{i-1} and D_i we can recover the second one. Eventually, we are able to compute $P(D_{N_1})$.

Now we have a more complicated diagram with two distinguished adjacent regions; we can then iterate the process: choose another edge of the first region[24] and make N_2 identical Ω_1 moves, with $N_1 \gg N_2 \gg 0$, in such a way that the curls are contained in the first region. Again, the first step is well defined, since such Ω_1^+ moves are the only ones that reach diagrams whose positive valence is increased by approximately $8N_1$ and do not increase the multiplicity of the multiedge. From the second step on, the lack of periodicity (see Remark 3.6) ensures that making a curl close to the previous one is the only way to create an Ω_1^+ multiedge. We still can recover $P(D_i)$ at every step.

We can fill up every edge of the region which once had a edges in the same fashion, and then move to another region. If at each step we start making curls on an edge bounding two regions whose number of edges is different enough,[25] and if we keep track of the number of curls added, we are sure that the moves are well defined on the graph and that we can compute the n–tuples.

Notice that we need to choose the numbers N_i incredibly big and suitably distant, with $N_i \gg N_{i+1}$ in order to avoid confusion and ultimately get a diagram $\widetilde{D}$ such that it has region sequence of the form

$$P(\widetilde{D}) = (N, 0, \dots, 0, 1, 0 \dots, 0, 1, 0, \dots),$$

where $N = \sum_i N_i$, and the minimal gap between two nonzero entries is $\gg 0$. Call a diagram with these properties *sparse*.

If the number of edges of the various regions are sufficiently spaced, then the previous claim applies,[26] and we can explicitly see which regions are adjacent to one another. However thus far we have only determined the dual graph to the knot projection as an *abstract graph*; we need a bit more work to find out the specific planar embedding of the dual, in order to get back the projection of D_0. It is well known that an abstract finite planar graph G, together with a *rotational system*, uniquely determines an embedding

[24] Such that it is not on the top of the Ω_1^+ moves we just made.

[25] This can be achieved by moving through adjacent regions. Notice that every time an edge is filled with N_i curls, we can "remember" the number of edges of the regions involved, and the number of curls made.

[26] Even if $p_1(\widetilde{D}) \neq 0$, the previous claim applies, thanks to the sparseness of $\widetilde{D}$ and to the fact that the only 1–regions in $\widetilde{D}$ are the ones created by the construction.

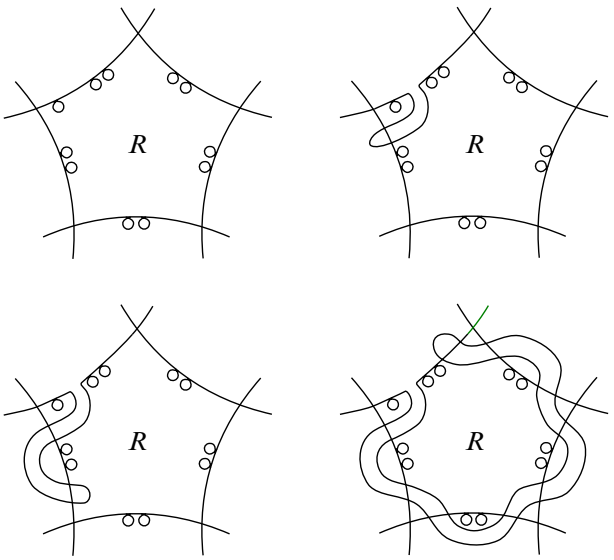

Figure 47: The choice of a local orientation system for the dual. By the previous construction the curls on the arcs can be on either side of an edge.

of G and thus G^*, which is the diagram projection we want. A local rotational system for a vertex $v \in G$ is just a choice of a cyclic order for the edges emanating from v. A rotational system for G is such a choice for each $v \in G$, and it is said to be *coherent* if all the local systems are coherently oriented.[27]

Choose a region R in the sparse diagram $\widetilde{D}$ and suppose that R is bounded by r edges. Choose an Ω_2^+ move that creates a new bigon and a new 4–region and increases the number of edges of two regions adjacent to R by 2; one example is shown in the top-right part of Figure 47. Notice that this choice is well defined on the graph thanks to the sparseness of $\widetilde{D}$. Indeed, we know the numbers r_i of edges of all the regions adjacent to R, and the Ω_2^+ moves of that kind are the only ones that change the positive valence by a value of $56 + 4r_j + 4r_i - 4r$ (as in (3-16)).

Then there are only two choices for a second Ω_2^+ move that creates a new bigon adjacent to R leaving all the 1–regions on the edge on the left of the newly created bigon, eliminates one bigon adjacent to one of the two regions whose vertices increased in the previous move (let us call this region M), and replaces it with a quadrilateral. Again, thanks to the sparseness of $\widetilde{D}$ we can identify such a move on the graph, by considering

[27] That is, given any two adjacent regions in the (embedded) dual of G, the orientations induced on the common edges do not coincide, as in Figure 48.

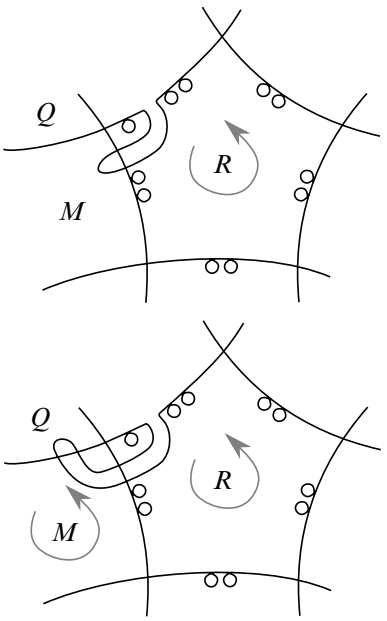

Figure 48: How to choose a coherent cyclic ordering for the orientation system.

the valence of the diagram reached. In fact, we can identify all the Ω_2 moves creating a new bigon adjacent to R and eliminating the other bigon, since these are the only ones changing the valence by[28] $40 + 4r - 2n_1 r_j + 2(n_1 + 2)(n_1 - 2)$ (as in (3-16)), where n_1 is the number of 1–regions on the left of the newly created bigon. In order to detect the correct Ω_2 moves it is sufficient to choose the one minimising the coefficient of r_j in the previous expression.[29] These two options correspond to the possible choices of over/under passing for the first Ω_2^+ move in Figure 47. These moves might also decrease by a lot the number of edges of M (according to how many curls are contained on the edge between R and M). Now we can repeat the process, following[30] Figure 47, until we get back to the first region that had its edges increased by 2 in the first move (but was not M). Keeping track of the various regions encountered during this process allows to reconstruct a local orientation system about the vertex corresponding to R in the dual graph. Since we know the numbers r_i of edges of all

[28] The following expression is valid if the 1–regions are inside R. To consider the other case it is enough to replace every n_1 in the expression with $2n_1$.

[29] We can do that thanks to the sparseness of the diagram.

[30] We only need to follow moves that do not *separate* curls lying on the same edge in different regions, and we can do that again using the difference in the valence together with the sparseness of the initial diagram.

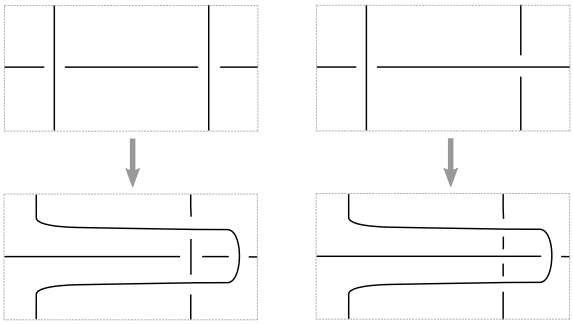

Figure 49: The two paths in $\mathcal{G}(K)$ and $\mathcal{G}(K')$. The grey arrows denote the sequence of $\Omega_2^+ - \Omega_3$ moves connecting the two diagrams.

the regions adjacent to R, and thanks to the sparseness of the diagram, this construction works even if two distinct edges of R are shared with the same region.

Again, this sequence of Ω_2^+ is only well defined up to a choice of over/under passing at each step, but this indeterminacy does not affect the result.

Finally, in order to get a proper orientation system for the dual, we need to be able to have a coherent way of orienting these local rotational systems we obtained. The process is shown in Figure 48. Once we have made the first Ω_2^+ move of Figure 47 (and thus chosen a clockwise or counterclockwise orientation), there are only two other Ω_2^+ moves that increase the number of regions with three sides by one and change the valence by approximately[31] $4q-2m$. This move will also increase by 2 the number of edges of the region denoted by Q; we are going to choose the only cyclic orientation based at the vertex in the dual, corresponding to the region M, that has Q after R. Repeating this process for all regions produces a well-defined and coherent orientation system for each vertex in the dual graph and hence uniquely determines the embedding of the dual and consequently the knot projection.

Now suppose we have a periodic knot type K; in order to repeat the previous strategy we need to be able, for each $D \in \mathcal{D}(K)$ with $p_1(D) = 0$, to find in a controlled way a sparse diagram $D' \in \mathcal{D}(K)$ and a sufficiently large ball, centred in D', such that all diagrams in this ball are nonperiodic. Since we can verify whether a diagram is periodic,[32] by changing the order of the Ω_1^+ multiedges appearing in the previous construction, and/or their valence, we can achieve a sparse configuration with these properties. □

[31] Lower case letters denote the number of edges in the corresponding regions.

[32] As explained in the proof of Theorem 3.23.

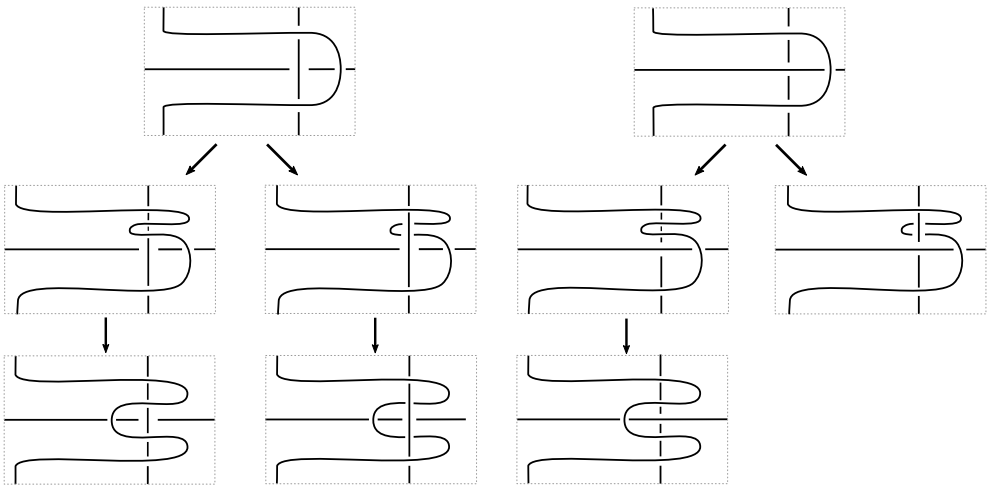

Figure 50: The two paths in $\mathcal{G}(K)$ and $\mathcal{G}(K')$. The top/middle arrows denote an Ω_2^+ and an Ω_3 move respectively.

To conclude the proof, we only need to be able to say that the only possible knots sharing all projections without curls are mirrors of one another.

Proposition 5.3 *The S^2–graph $\mathcal{G}(K)$ detects some diagrams of K up to mirroring.*

Proof Suppose we have two knots K and K' sharing the same graph. Take a vertex D of $\mathcal{G}(K)$ such that $\#\Omega_1^-(D) = 0$. The corresponding vertex D' in the isomorphic graph $\mathcal{G}(K')$ will have the same knot projection as D by Proposition 5.2. Hence if $K \neq K'$ the diagrams must differ in at least one crossing. If they differ in *all* crossings, then K' is the mirror of K, and we are done. Otherwise there must be a pair of crossings that up to mirroring looks like the pair in the top part of Figure 49 in D and D'.

Now, perform the sequence of Ω_2^+–Ω_3 moves that takes the upper diagrams in Figure 49 and ends in the lower ones. Note that these paths are well defined, since all the diagrams involved respect the condition $p_1 = 0$; thus we are able to actually determine the effect of these moves on the projections by Proposition 5.2. Consider Figure 50; there are two distinct sequences of Ω_2^+–Ω_3 (differing by the choice of over/under passing for the first Ω_2^+ move) starting from the diagram on the top left of Figure 50 and ending in two different diagrams sharing the same projection. On the other hand, there is only one way to perform an Ω_2^+ from the diagram on the top right of Figure 50 in order to be able to complete the sequence with an Ω_3 and obtain a diagram with the same projection as the other two. Again, these moves are all well defined thanks to Proposition 5.2.

Hence the two graphs cannot coincide, since there is a path in one of the graphs which is not present in the other one, and we can conclude. □

The three previous propositions together with Proposition 3.11 can be easily seen to imply Theorem 1.1, but as a matter of fact the result proved is even stronger, since it allows to recover the actual diagrams[33] represented by some specific vertex of the graph, and not only its knot type. From the proof of Proposition 5.2 we are actually obtaining an embedding for the graph which is dual to the knot projection corresponding to the diagram D_0. Hence, this proves that we can actually get back the shape of any diagram not containing any region with one edge (in the nonperiodic case).

This next result follows directly from the proofs of Propositions 5.1–5.3:

Corollary 5.4 *Let K be a knot. For every vertex $D \in \mathcal{G}(K)$ there exists an integer $R > 0$ such that $S_R(D)$ is* ***characterising****, meaning that this graph can only appear in $\mathcal{G}(K)$. Moreover, in the nonperiodic case, R is computable.*

A similar argument should guarantee the completeness of the planar $\mathcal{R}$–graphs, even though the whole process is complicated by the fact that the presence of the external region does not allow a straightforward adjustment of Proposition 5.1.

6 The blown-up Gordian graph

We can unify the Gordian and Reidemeister graphs in a single object, by a sort of "blowup" construction; just replace each vertex of the Gordian graph with the corresponding $\mathcal{G}(K)$. The edges between two knots in the Gordian graph can be split into edges between the diagrams realising the crossing changes.

Definition 6.1 Define the *blown-up Gordian graph* $\mathcal{G}^*_P$ (respectively $\mathcal{G}^*$) as the graph whose vertices are knot diagrams in the plane (respectively in S^2) up to the corresponding notion of diagram isotopy; there is an edge between two vertices if and only if they are connected by a single Reidemeister move or a crossing change.

As in the previous setting, the valence of each vertex is finite. For nonperiodic diagrams,

$$v_*(D) = v(D) + \mathrm{cr}(D),$$

where $v_*(D)$ denotes the valence of D in $\mathcal{G}^*_P$, and $v(D)$ is the valence of D in the corresponding $\mathcal{R}$–graph. For periodic diagrams we only get an inequality.

[33] For knots $K \neq \bigcirc$.

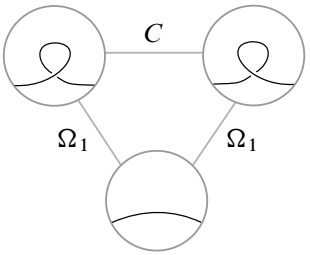

Figure 51: The only other possible length 3 path in $\mathcal{G}^*$ which is not a path in any $\mathcal{G}(K)$.

Note that $\mathcal{G}^*$ admits an order 2 automorphism, induced by changing all crossings of each diagram, ie taking the mirror image. The only fixed points of this automorphism are the diagrams of amphicheiral knots which are equivalent to their mirror up to planar isotopy.

Remark 6.2 There are embeddings of $\mathcal{G}_P(K) \hookrightarrow \mathcal{G}_P^*$ and $\mathcal{G}(K) \hookrightarrow \mathcal{G}^*$ for each $K \in \mathcal{K}$, and there are many crossing-change edges in both $\mathcal{G}_P^*$ and $\mathcal{G}^*$ connecting two diagrams in the same isotopy class; according to the *cosmetic crossing conjecture* [13, Problem 1.58] all these should correspond to nugatory crossings. It would be interesting to explore the possible applications of these graphs to the conjecture.

If we look at the ball of radius 1 in $\mathcal{G}^*$ about a diagram D, we find all the length 3 paths of Theorem 3.2, together with a new configuration, shown in Figure 51. The fact that the only new triangle[34] appearing is actually this one follows easily by considering the Arnold and Hass–Nowik invariants, together with crossing number and writhe, as in Theorem 3.2. More precisely, using these invariants we can restrict to cycles of the form $\Omega_1 - \Omega_1 - C$, where the Ω_1 moves create crossings of opposite sign, and C denotes a crossing change. Then, with the same line of thought as in Theorem 3.2, we can prove that the curls must lie in the same region by taking in account the self-touching number. It follows that the regions under the two curls have the same number of edges. However, we are not able to prove that the crossing change happens exactly on the curls.

By extending the proof of Theorem 3.23 to the blown-up graph it is possible to prove the analogous result; namely that $\mathcal{G}^*$ detects the crossing changes and Reidemeister moves.

It is also possible to employ the blown-up graph in quite different contexts, like modelling DNA pathways or considering walks on it to produce cryptographic protocols. These applications will be the subject of upcoming works.

[34] That is, triangles that contain at least one crossing change.

Other related ongoing projects include a translation of the concepts outlined in this paper to a plethora of other settings; the rough idea is the following: given a recipe to present knots, and a finite set of moves to pass between equivalent presentations, one obtains a related graph. In an upcoming paper we are going to study what happens in the case of grids, braids, tangles and pointed and framed diagrams. We will also provide computations for the corresponding diagram complexity invariants, for low-crossing knots and some infinite family. Moreover, we are going to explore some of the connections between the $\mathcal{R}$–graphs defined here and the topology of the discriminant hypersurfaces in Arnold's and Vassiliev's constructions [2; 18].

We conclude with some questions.

Question 6.3 *Does there exists a periodic knot type K such that the minimum of the complexity (in either graph) is not attained at a periodic diagram*?

Question 6.4 *Are all knot types simple* (as in Definition 3.8)?

Question 6.5 *If a diagram $D \subset S^2$ of a nontrivial knot is periodic, is it true that all other diagrams in $S(D)$ are not periodic*?

Question 6.6 *Is the S^2–graph obtained from a minimal set of Reidemeister moves a complete invariant*?

Question 6.7 *To what extent do the filtrations $\mathcal{F}(K)$ and $\widetilde{\mathcal{F}}(K)$ classify knot types*?

The following (hard) question was suggested by Lackenby:

Question 6.8 *We have shown that the isomorphism class of the S^2–Reidemeister graph is a complete knot invariant. Is the* ***quasi-isometry*** *class of the graph a complete invariant, or if not, to what extent does it distinguish inequivalent knots or detect interesting properties of the knot*?

The following question was asked by D Cimasoni:

Question 6.9 *Which knot invariants* (such as genus, absolute value of the writhe, polynomials, ...) *can be extracted from the $\mathcal{R}$–graphs*?

Question 6.10 *Does the connectivity of the $\mathcal{R}$–graphs coincide with the minimal complexity*?

Question 6.11 *We have exhibited an infinite set of spheres representing nontrivial elements of $H_2(\mathcal{CG}(K);\mathbb{Z})$* (in Section 4). *Are there any embedded closed surfaces of higher genus*? (Note that such surfaces would automatically be nontrivial in homology.)

References

[1] **J W Alexander**, **G B Briggs**, *On types of knotted curves*, Ann. of Math. 28 (1926/27) 562–586 MR

[2] **V I Arnold**, *Plane curves, their invariants, perestroikas and classifications*, from "Singularities and bifurcations" (V I Arnold, editor), Adv. Soviet Math. 21, Amer. Math. Soc., Providence, RI (1994) 33–91 MR

[3] **S Baader**, *Note on crossing changes*, Q. J. Math. 57 (2006) 139–142 MR

[4] **R Budney**, *Structure of Gordian graph of knots*, MathOverflow post (2016) Available at `https://mathoverflow.net/q/255140`

[5] **A Coward**, **M Lackenby**, *An upper bound on Reidemeister moves*, Amer. J. Math. 136 (2014) 1023–1066 MR

[6] **T Ekholm**, **L Ng**, **V Shende**, *A complete knot invariant from contact homology*, Invent. Math. 211 (2018) 1149–1200 MR

[7] **E Flapan**, *Infinitely periodic knots*, Canad. J. Math. 37 (1985) 17–28 MR

[8] **C M Gordon**, **J Luecke**, *Knots are determined by their complements*, J. Amer. Math. Soc. 2 (1989) 371–415 MR

[9] **R Halin**, *Über die Maximalzahl fremder unendlicher Wege in Graphen*, Math. Nachr. 30 (1965) 63–85 MR

[10] **J Hass**, **T Nowik**, *Invariants of knot diagrams*, Math. Ann. 342 (2008) 125–137 MR

[11] **D Joyce**, *A classifying invariant of knots, the knot quandle*, J. Pure Appl. Algebra 23 (1982) 37–65 MR

[12] **L H Kauffman**, **S Lambropoulou**, *Hard unknots and collapsing tangles*, from "Introductory lectures on knot theory" (L H Kauffman, S Lambropoulou, S Jablan, J H Przytycki, editors), Ser. Knots Everything 46, World Sci., Hackensack, NJ (2012) 187–247 MR

[13] **R Kirby**, *Problems in low dimensional manifold theory*, from "Algebraic and geometric topology, II" (R J Milgram, editor), Proc. Sympos. Pure Math. XXXII, Amer. Math. Soc., Providence, RI (1978) 273–312 MR

[14] **J Marché**, *Comportement à l'infini du graphe gordien des nœuds*, C. R. Math. Acad. Sci. Paris 340 (2005) 363–368 MR

[15] **Y Miyazawa**, *A distance for diagrams of a knot*, Topology Appl. 159 (2012) 1122–1131 MR

[16] **M Polyak**, *Minimal generating sets of Reidemeister moves*, Quantum Topol. 1 (2010) 399–411 MR

[17] **K Reidemeister**, *Elementare Begründung der Knotentheorie*, Abh. Math. Sem. Univ. Hamburg 5 (1927) 24–32 MR

[18] **V A Vassiliev**, *Cohomology of knot spaces*, from "Theory of singularities and its applications" (V I Arnold, editor), Adv. Soviet Math. 1, Amer. Math. Soc., Providence, RI (1990) 23–69 MR

Mathematical Institute, University of Oxford
Oxford, United Kingdom

Mathematical Institute, University of Oxford
Oxford, United Kingdom

barbensi@maths.ox.ac.uk, celoria@maths.ox.ac.uk

Received: 24 January 2018 Revised: 7 January 2019