

Unknotted Cycles

Christopher Cornwell Nathan McNew

Department of Mathematics
Towson University
Maryland, U.S.A.

{ccornwell,nmcnew}@towson.edu

Submitted: Feb 1, 2022; Accepted: Jul 19, 2022; Published: Sep 9, 2022

© The authors. Released under the CC BY-ND license (International 4.0).

Abstract

Noting that cycle diagrams of permutations visually resemble grid diagrams used to depict knots and links in topology, we consider the knot (or link) obtained from the cycle diagram of a permutation. We show that the permutations which correspond in this way to an unknot are enumerated by the Schröder numbers, and also enumerate the permutations corresponding to an unlink. The proof uses Bennequin's inequality.

Mathematics Subject Classifications: 05A15, 57K10

1 Introduction

A convenient way to visualize a permutation is to draw a plot of the permutation on an $n \times n$ grid, placing a dot in the box at each of the locations $(i, \sigma(i))$. The permutation's cycle structure can be represented by making the plot to be a *cycle diagram* [13]. At each index i we draw a vertical line from (i, i) to the point $(i, \sigma(i))$, followed by a horizontal line to $(\sigma(i), \sigma(i))$. If i is a fixed point, $i = \sigma(i)$, no additional lines are drawn. The result is a diagram in which the cycles of the permutation can be traced out along the lines of the diagram in a natural way.

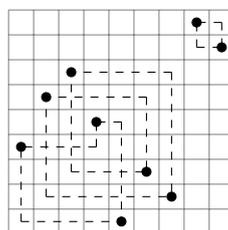


Figure 1: The cycle diagram of $\pi = 467513298$.

For example, from the cycle diagram for the permutation $\pi = 467513298$ (written in one line notation) depicted in Figure 1, one can readily identify the cycle decomposition $\pi = (145)(2637)(89)$ (written in cycle notation) by tracing out the lines of the diagram.

Note that the only corners in a cycle diagram occur at the plotted points of the permutation and along the line $y=x$. The appearance of these diagrams strongly resembles grid diagrams which are a useful tool in the study of knots in topology.

Formally, a grid diagram is an $n \times n$ lattice where each row, and each column, has exactly two marked boxes (traditionally marked X and O) and a line is drawn between the marked boxes in each row and column. The diagram is interpreted as a knot (an embedding of S^1 in \mathbb{R}^3) or a link (multiple copies of S^1) by designating all of the vertical lines to be overcrossing and the horizontal lines as undercrossing. (Section 2 provides an overview of necessary basic notions from knot theory, see also [8, 11, 20].)

Crucially, the only distinctions between diagrams that are valid cycle diagrams and those that are valid grid diagrams are:

- In cycle diagrams one of the two designated points in each row/column must lie on the $y = x$ line, which isn't necessarily the case for grid diagrams.
- In grid diagrams, it is not allowed to have a single point in a row or column as occur in cycle diagrams when there are fixed points.

In light of the second point, so long as we take a permutation without fixed points (called a *derangement*) then by drawing its cycle diagram and interpreting it as the grid diagram of a link we can build a link corresponding to any derangement. We refer to this link as the *link associated to a permutation* or, when the permutation is a cycle, as the *knot associated to a cycle*. We do not associate a link to a permutation that is not a derangement.

Many natural questions arise, including:

1. Which knots (links) are associated to some derangement?
2. How many different derangements are associated with a given link?

While we do not have a complete answer to question (1), the answer is certainly not “all knots” (or links), as explained at the end of Section 2.1.

As to question (2), in this paper we enumerate the cycles that are associated to the unknot (called an *unknotted cycle*) as well as permutations associated to an unlink (*unlinked permutations*). Our main result is the following.

Theorem 1. *The unknotted cycles are enumerated by the shifted sequence of (large) Schröder numbers S_n . Enumerating these numbers as $S_1 = 1, S_2 = 2, \dots$. The number of unknotted cycles of length n is S_{n-1} .*

The Schröder numbers count a wide array of combinatorial objects, notably separable permutations of length n , and lattice paths from $(1,1)$ to (n,n) consisting of north $(0,1)$, east $(1,0)$ and northeast $(1,1)$ steps which never go above the diagonal [21]. Their

generating function is $S(x) = \sum_{n=1}^{\infty} S_n x^n = \frac{1}{2} (1 - x - \sqrt{1 - 6x + x^2})$, which satisfies the recursion $S(x) = x + xS(x) + S(x)^2$. Asymptotically [18, Ex. 2.2.1-12]

$$S_n \sim \frac{\sqrt{2} - 1}{2^{3/4} \sqrt{\pi}} (3 + \sqrt{8})^n n^{-3/2}. \quad (1)$$

The relationship between cycle diagrams and grid diagrams does not appear to have been considered before in the literature. A related body of work is the study of random knots, particularly via the *random grid model* [15]. A common question in this area considers the probability of a random knot being equivalent to a fixed knot K . Particularly, for each given integer $n > 0$, random knot models select a knot K_n ; as $n \rightarrow \infty$, does the probability that K_n is equivalent to K approach zero, and what are the asymptotics? A broadly construed conjecture is the Frisch-Wasserman-Delbrück conjecture (see [15] for discussion).

In the *random grid model*, K_n is given by selecting a random pair of n -cycles (σ, τ) , independently and uniformly. Let σ be $(\sigma_1, \dots, \sigma_n)$ in cycle notation (so $\sigma(\sigma_i) = \sigma_{i+1}$ for each i , $\sigma_{n+1} = \sigma_1 = 1$, say), and likewise for τ . A grid diagram is then determined by drawing, for each $1 \leq i \leq n$, a vertical line from (σ_i, τ_i) to (σ_i, τ_{i+1}) , and from there a horizontal line to $(\sigma_{i+1}, \tau_{i+1})$. Then K_n is the knot of this grid diagram (a knot since σ and τ are n -cycles). In this model, Witte has shown that the probability of K_n being a given knot (e.g. the unknot), is $\mathcal{O}(n^{-1/10})$ as $n \rightarrow \infty$ (this is implied by Theorem 6.0.1 of [25]).

When $\sigma = \tau$, the grid diagram in the above model is the cycle diagram of σ . So, on the diagonal of the random grid model, Theorem 1 gives an exact probability of K_n being the unknot, equal to $\frac{S_{n-1}}{(n-1)!} \sim \frac{5\sqrt{2}-7}{2^{5/4}\pi n} \left(\frac{e(3+\sqrt{8})}{n}\right)^n$ using (1) and Stirling's formula. Hence, the probability of K_n being the unknot decays super-exponentially as $n \rightarrow \infty$.

We show that unknotted cycles are counted by the Schröder numbers by establishing a bijection between them and the rooted-signed-binary trees defined in Section 3. We first define a way to construct an unknotted cycle of size $n + 1$ from a rooted-signed-binary tree of size n in Section 4, then show in Section 4 that it is well defined and one-to-one on equivalency classes of these trees. Finally, in Section 5 we use results from topology to prove that it is surjective – all unknotted cycles can be obtained in this way. Finally, we enumerate the unlinked permutations in Section 6 and note a potential relationship to the Diaconis-Graham inequality.

2 Knots and Links

Informally, in this article, a *knot* is a closed curve in \mathbb{R}^3 which has no self-intersections. In addition, two knots K and K' are considered to be *equivalent* if there is a continuous deformation that takes K to K' , such that, at each time during the deformation of K , it is a knot. More precisely, a knot K is the image of a piecewise \mathcal{C}^1 -embedding, $S^1 \hookrightarrow \mathbb{R}^3$, and K is equivalent to a knot K' if there is an ambient isotopy of \mathbb{R}^3 (see [19]) carrying one to the other. When there are multiple knots, no two of which intersect, the multi-curve

is called a *link*. The notion of equivalent links is analogous to equivalent knots. We use “links” inclusively, so a link could be a knot (having just one component curve).

A link is *oriented* by choosing a consistent positive tangent along each component curve; in figures the positive tangent will be indicated by an arrow.



Figure 2: Positive and negative crossings.

2.1 Link diagrams

Often, knots and links are studied via “sufficiently generic” projections to a plane, in which any self-intersections that arise from the projection have independent tangent directions. The self-intersections are *crossings*. Projections of equivalent links can have different numbers of crossings. The projection plane is understood to have a positive normal direction, allowing us to say that one of the branches at the crossing has a “higher” projection preimage; this branch is the *overcrossing strand*, and the other is the *undercrossing strand*. In figures, the overcrossing strand appears to pass on top of the other. With this crossing information, the projection is called a *knot (or link) diagram*. A given link diagram determines the corresponding link, up to equivalence.

In the diagram of an oriented link, a sign is given to each crossing. If, possibly after a rotation, a neighborhood of the crossing appears as on the left side of Figure 2 then the crossing is positive (note the directionality of the arrows). Otherwise, it appears as on the right of the figure and the crossing is negative. The *writhe* of the diagram equals the number of positive crossings minus the number of negative crossings. The diagram of Figure 3 has writhe equal to 0.

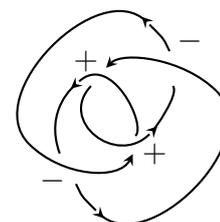


Figure 3: An oriented knot diagram which has writhe equal to 0.

Given a link K , we will regularly discuss a related link that is known as the *mirror (image) of K* . Owing to K being a subset of \mathbb{R}^3 , define the mirror of K to be a link equivalent to the image of K under the orientation-reversing map $(x, y, z) \mapsto (x, y, -z)$ sending \mathbb{R}^3 to itself. We may use a link that is equivalent to K instead; the mirror of a link is well-defined on equivalence classes. Additionally, we may understand the mirror through link diagrams. Given a link diagram of K , at every crossing of the diagram change which strand is overcrossing and which is undercrossing. The result is a link diagram for the mirror of K (for more details, refer to [9, Chapter I.4]).

The cycle diagrams introduced in Section 1 are examples of link diagrams; each crossing has a vertical overcrossing strand and a horizontal undercrossing strand. The choice to set vertical strands to be overcrossing is standard when using grid diagrams in the knot theory literature. A consequence of this choice is that the knots associated to cycle

diagrams are mirror images of knots from a well-studied class, *positive knots*, rather than being in that class (see more discussion below). Nevertheless, we find it preferable to follow the literature on grid diagrams.

As a convention, the orientation of the link associated to a cycle follows the order in which the permutation goes through the indices. That is, vertical (resp. horizontal) segments above the diagonal are oriented up (resp. rightward), and below the diagonal are oriented down (resp. leftward).

Recalling the sign given to crossings in an oriented link diagram, as in Figure 2, note that every crossing that appears in a cycle diagram is a negative crossing. Hence, if a knot is associated to a cycle diagram, its mirror is in a class of knots called *positive knots*. Many knot types, even among those with a small number of crossings, do not fit into this class. Some obstructions to being positive are known (e.g. see [7]).

A well-known construction from knot theory is the *closure of a braid*, which associates a link to an element of the braid group B_n . For every link there exists $n \geq 1$ so that the link is the closure of some braid in B_n (for a survey, see [4]). The most common presentation of B_n uses the *Artin generators*. If we restrict to braids expressible with only positive powers of Artin generators, the associated links are *positive braid closures*, a subclass of positive links with interesting applications in geometric topology.

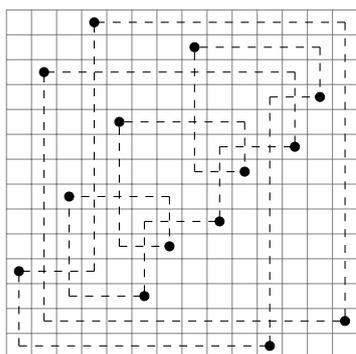


Figure 4: A cycle associated to the mirror of $14n_{5644}$, which is not a positive braid closure.

From the form of the cycle diagram construction, it is tempting to think that knots associated to cycles are mirrors of positive braid closures, and this is the case for all small cycles. However, it is not true in general. Consider the knot associated to $(4, 12, 7, 14, 10, 3, 5, 13, 6, 8, 1, 9, 11, 2)$, a cycle in one-line notation, depicted in Figure 4. Call the mirror of this knot K_0 , which is a prime knot. The grid diagram for this cycle has 14 crossings, so K_0 is among knots with crossing number 14 or less (which have been enumerated). The only such knot with a Jones polynomial in agreement with K_0 is $14n_{5644}$ (in the Hoste-Thistlethwaite enumeration). We conclude that K_0 is the knot $14n_{5644}$, which is not a positive braid closure (cf. [1]).¹

2.2 Boundaries of surfaces

Every (oriented) link can be realized as the boundary of a connected, oriented surface in \mathbb{R}^3 . Such a surface is called a *Seifert surface* of the link, after Herbert Seifert who gave

¹This cycle was found by finding the knot associated to each cycle of length at most 14, computing its Jones polynomial and comparing it to the Jones polynomials of all positive braid closures of genus at most 5 enumerated in [1]. The Jones polynomial of (the mirror of) every knot, associated to a cycle of length at most 13, having minimal degree 5 or smaller, agrees with the Jones polynomial of a positive braid closure, so the counterexample given is minimal in that sense. The knot $14n_{5644}$ has genus 5 and is listed in [1] among the positive fibred arborescent knots of weight 2, but not the positive braid closures.

an algorithm for producing one, given a link diagram [23] (the existence of a surface was known earlier [16]).

If a link L has b component curves then a Seifert surface Σ of L will have b boundary components. A closed orientable surface $\widehat{\Sigma}$ may be defined as an identification space from Σ (starting with the disjoint union of Σ and b disjoint disks, identify the boundary of each disk to one of the boundary components of Σ). The *genus* of Σ is the genus of $\widehat{\Sigma}$. If g is that genus, then $\chi(\widehat{\Sigma}) = 2 - 2g$; it must be that the Euler characteristic of Σ is $2 - 2g - b$. The minimal genus of a Seifert surface of a link L is called the *genus* of L , written $g(L)$. Equivalent links have equal genus. In addition, the mirror of L has genus equal to the genus of L , since the map $(x, y, z) \mapsto (x, y, -z)$ carries any orientable surface to another orientable surface with the same genus.

A part of Seifert’s algorithm, to be used in Section 5, involves determining Seifert circles from an oriented link diagram. To define Seifert circles, consider a crossing (with an orientation, so each strand has an incoming and outgoing end). Remove the crossing point and connect the coherently oriented strands (“smooth” the crossing), as in Figure 5. By smoothing every crossing of the diagram, we obtain a collection of pairwise disjoint circles in the plane (circles in a topological sense). These are the *Seifert circles* of the diagram. See an example in Figure 12.

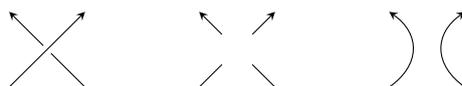


Figure 5: Smoothing a crossing: remove the crossing, then reconnect the remaining edges in the same orientation without crossing.

2.3 Legendrian links

Related to the modern study of contact geometry are *Legendrian links*. Loosely speaking, given a contact structure on \mathbb{R}^3 , a knot or link is *Legendrian* if it satisfies a tangency condition based on the contact structure (additionally, equivalence of two Legendrians is constrained by the tangency condition). We refer the interested reader to the survey article [14].

In the “standard” contact structure, Legendrian links are often studied via a specific planar projection, the *front projection*. This projection has some peculiarities: it has no vertical tangencies, but has *cusps* (locally like the cuspidal cubic); also, at a crossing, the more negatively-sloped branch is always the overcrossing strand. A link diagram having these two properties (and being smooth except at cusps), is the front projection of a (unique) Legendrian link.

A grid diagram determines a knot or link (simply viewing it as a link diagram, with vertical overcrossings). We will also determine a Legendrian link from a grid diagram as follows.² First, rotate the grid diagram 45° clockwise. Then, turn what were origi-

²Our method of determining a Legendrian from a grid diagram is not typical in the literature [20]; however, our method is closely related, and convenient for our purposes.

nally lower-left and upper-right corners into cusps, and smooth out the upper-left and lower-right corners (now local extrema vertically). Finally, at each crossing interchange overcrossings with undercrossings (see Figure 6; the grid diagram has vertical overcrossing strands, and the front projection has negatively-sloped overcrossing strands).

For a grid diagram D , thinking of D as a link diagram, let $K(D)$ be the associated link. Let $\Lambda(D)$ be the Legendrian link of the front projection that we described above. Then, as a regular link, $\Lambda(D)$ is equivalent to the mirror of $K(D)$. This fact will have very little effect on our arguments. Indeed, we use the construction when $K(D)$ is an unknot; an unknot and its mirror image are equivalent knots. As crossings of a cycle diagram D are negative, every crossing of $\Lambda(D)$ is positive for us.

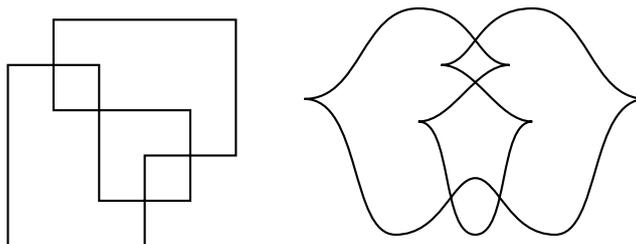


Figure 6: A grid diagram (left) and its front projection (right).

An invariant of Legendrian knots and (oriented) links that is of interest to us is the Thurston-Bennequin number. Given a Legendrian Λ , let F_Λ be its front projection. The *Thurston-Bennequin* number $\text{tb}(\Lambda)$ is equal to the writhe of F_Λ minus one-half the number of cusps. Note that one-half the number of cusps is the same as the number of “right cusps” – cusps that point to the right:

$$\text{tb}(\Lambda) = \text{writhe}(F_\Lambda) - (\#\text{right cusps in } F_\Lambda).$$

In Section 5 we will use the Bennequin-Eliashberg inequality [12], which says that if Λ is equivalent, as a regular knot or link, to some K , then

$$\text{tb}(\Lambda) \leq 2g(K) - 1.$$

3 Signed Trees

Various authors [5, 24] introduce *separating trees* to study separable permutations. A separating tree is a rooted binary tree in which each internal node is designated as either positive or negative, and then trees are divided into equivalence classes under certain allowable tree rotation operations. (See also [6].) Using separating trees as motivation, we define a similar structure, which will be useful in the enumeration of unknotted cycles.

Definition 2. A **rooted-signed-binary tree** is a rooted binary tree in which each node except the root is given a sign, positive or negative. Furthermore, we say two binary rooted trees are equivalent if one can be obtained from another by a series of tree rotations. A tree rotation (see Figure 7 for an example) is allowed at a given node if either:

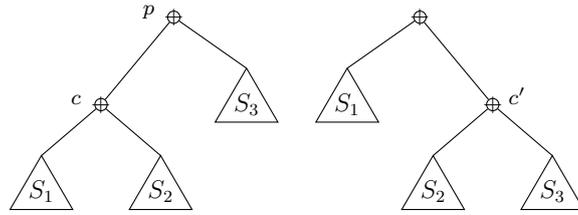


Figure 7: The result of a single tree rotation. The child node c is described as being *rotated into place* of its parent p . The node c' is *created* by the rotation. As represented by the triangles, descendants to the right of p become descendants to the right of c' ; descendants to right of c become descendants to the left of c' . The relative positions internal to the triangles S_i remain unchanged.

1. The parent node has the same sign as the child rotating into its place.
2. The parent node is the root, in which case the newly created node is assigned the same sign as the node that was rotated into the position of the root.

For example, the seven trees depicted in Figure 8 are all equivalent, and represent all of the allowed rotations of the given tree.

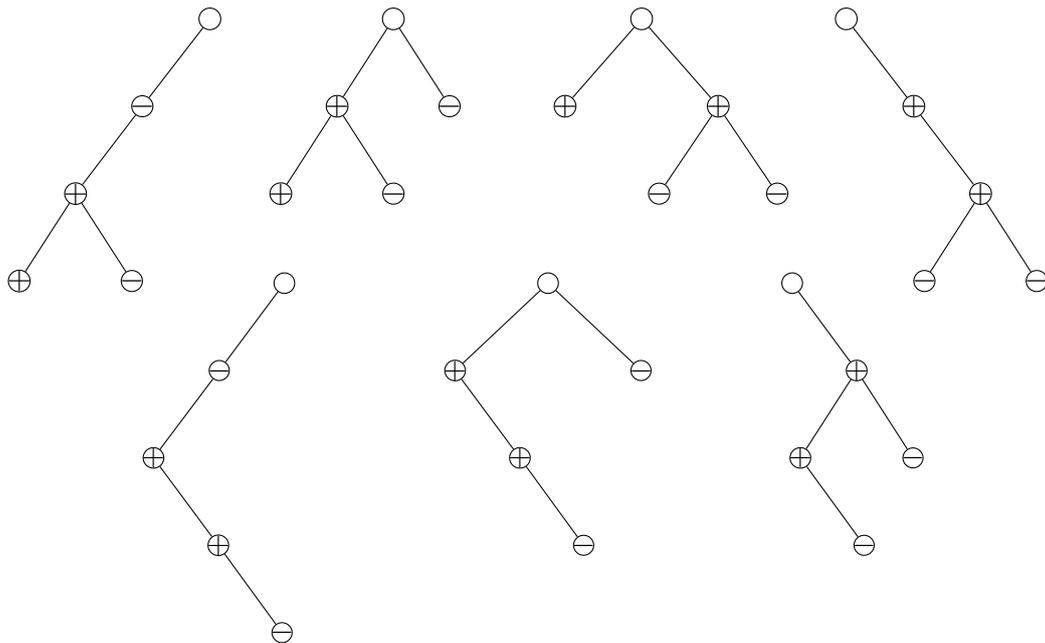


Figure 8: All seven unique tree rotations of a rooted-signed-binary tree. The top row depicts all of the possible trees obtained starting with the leftmost tree and performing rotations at the the root vertex. The second row shows the trees that can be obtained by rotating at other vertices.

Frequently in the remainder of the paper we will use the term rooted-signed-binary-trees to refer to equivalence classes of such trees. As our rooted-signed-binary-trees differ

slightly from the separating trees which are known to be in bijection with separable permutations, we give a proof that the number of such trees with n nodes is counted by the Schröder Numbers.

Proposition 3. *The number of (equivalence classes of) rooted-signed-binary-trees with n nodes is counted by the n -th Schröder number, S_n .*

Proof. We show this using generating functions. Let $T(x)$ be the ordinary generating function for such trees, and take as our representative for each equivalence class of rooted-signed-binary-trees the tree in which all possible left rotations have been made. Note that such a tree has a root node with no right child.

There are three possibilities for the left child: no left child, a left child which itself has no right child, or a left child whose right child has the opposite sign. These three cases are counted by x , $2xT(x)$, and $2xT(x) \left(\frac{T(x)}{2x} - \frac{1}{2} \right) = T(x)^2 - xT(x)$ respectively. The last two require brief explanations. Trees where the left child of the root has no right child can be constructed by taking any tree (counted by $T(x)$), assigning the root node either of 2 signs and then making that node the left child of a new signless root node, (increasing the size by 1) giving $2xT(x)$.

Now, if the left child of the root has a right child of the opposite sign, the possibilities for that right child can be counted by $\frac{T(x)/x}{2} - 1$. The $T(x)/x$ counts the trees with the unlabelled root node removed, dividing by 2 accounts for the fact that the sign of this right child must be the opposite of the node above it, and subtracting one eliminates the possibility of this being empty. Adding these three cases together and simplifying we obtain

$$T(x) = x + xT(x) + T(x)^2,$$

the same recurrence as the generating function of the Schröder numbers. □

4 The bijection

In this section we establish a bijection between (equivalence classes of) rooted-signed-binary-trees and the unknotted cycles. In doing so we will frequently build these trees by inserting (or removing) leaf nodes into existing trees. If T is a rooted-signed-binary tree and v is a (non-root) leaf of that tree, then we denote by $T - \{v\}$ the tree obtained by removing the node v . Sometimes we may also remove multiple nodes (and write $T - \{v, \dots\}$). However, we will never remove an internal node without also removing its descendants.

We enumerate the positions that a new leaf-node could be inserted from left to right and refer to them as follows. (Note that in a tree with n nodes there are always $n + 1$ positions where a new leaf can be inserted.)

Definition 4. Suppose a tree T has n nodes, and v is a leaf of T . We say that v is in **relative position** i , where $1 \leq i \leq n$, if there are exactly $i - 1$ places where a leaf could be inserted in T to the left of v (not counting the left child position of v). We also say

that a node is inserted in position i , meaning that after insertion the node is in relative position i .

Remark 5. The relative position of a leaf is unaffected by any tree rotation of the tree so long as it was not being rotated into the place of its parent.

We can now describe how to construct an unknotted cycle from a fixed rooted-signed-binary-tree. We will define our construction by describing how each node added to the tree affects the corresponding cycle. We start with the rooted-signed-binary-tree consisting only of the unlabelled root, which corresponds to the trivial cycle, 21 which is clearly unknotted.



Figure 9: The cycle 21.

At any point in this construction our tree will have as many places where a new node can be added as the length of the permutation thus far constructed (which is one more than the number of nodes in the tree).

When the tree consists of only the root node, there are two places where a node can be added, as either left or right children of the root node, corresponding to positions 1 and 2 of the trivial cycle respectively.

Now, when a new node is added a new point is inserted into the diagram, and other points are shifted up and to the right. We define functions

$$\xi_m(k) = \begin{cases} k & k < m \\ k + 1 & k \geq m. \end{cases} \quad (2)$$

When a node is added to the tree in relative position i , it changes the corresponding cycle $\sigma = s_1 s_2 \cdots s_n$ (in one line notation) into a new cycle $\sigma' = s'_1 s'_2 \cdots s'_n s'_{n+1}$ according to the following rules:

1. If a positive node is added, then σ changes as below.

$$s_1 s_2 \cdots s_n \rightarrow \xi_{i+1}(s_1) \xi_{i+1}(s_2) \cdots \xi_{i+1}(s_{i-1}) [i+1] \xi_{i+1}(s_i) \cdots \xi_{i+1}(s_n) =: \sigma'$$

To be clear, this means that $s'_j = \xi_{i+1}(s_j)$ for $1 \leq j \leq i - 1$, $s'_i = i + 1$, and that $s'_j = \xi_{i+1}(s_{j-1})$ for $i + 1 \leq j \leq n + 1$.

2. If a negative node is added, then σ changes as below.

$$s_1 s_2 \cdots s_n \rightarrow \xi_i(s_1) \xi_i(s_2) \cdots \xi_i(s_i) [i] \xi_i(s_{i+1}) \cdots \xi_i(s_n) =: \sigma'$$

Note that, in this case, $s'_{i+1} = i$.

In terms of the cycle diagrams, this has the effect of taking one of the corners where the diagram made a right angle at the $y = x$ line and changing it into a notch or a kink with a new off-diagonal point in the cycle diagram. The exact change depends on the behavior of the lines in the diagram prior to the insertion, and are summarized in Table 1.

Before insertion	After inserting \oplus	After inserting \ominus

Table 1: A graphical depiction of how the cycle diagram changes when a node is added to the tree. The first column depicts the corner on the diagonal before the new point is inserted. Note, when a \oplus node is inserted, a corner is created above the diagonal, and conversely, if a \ominus is inserted, the new corner is below the diagonal.

From these pictures it is clear visually that these changes will not affect the knot of the cycle diagram, up to equivalence. (Formally, each of the changes is either planar isotopy or a Reidemeister I move.) Thus if a given cycle is unknotted before one of these operations is performed, it will still correspond to the unknot afterward. The proof of the following proposition follows immediately.

Proposition 6. *Inserting an off-diagonal element to a permutation σ (an $i+1$ in position i or an i in position $i+1$) and shifting the points above or to the right in the cycle diagram results in a permutation associated to the same link as σ .*

By repeated application of this proposition we get the following theorem.

Theorem 7. *If a cycle σ is obtained from a rooted-signed-binary-tree by the construction above (processing nodes in some order) then σ is an unknotted cycle.*

We illustrate this by building the cycle corresponding to the tree in Figure 8. We use the second diagram of the top row depicted in that figure. The reader is invited to verify the same cycle is obtained for any equivalent tree and irrespective of the order the nodes of the tree are considered, as we subsequently prove.

Example 8. We consider the nodes from the second tree appearing in Figure 8 one at a time. Starting with the root node, we have the trivial cycle 21, shown in Figure 9.

We first process the positive, left child of the root. This positive node is in relative position 1, so we insert 2 into position 1, obtaining 231. We could also have found this from the cycle diagram, noting the corner in position 1 was a lower left corner (row 3 of table 1) and replacing the corner in the diagram with the picture in the second column. The cycle 231 is depicted first in figure 10.

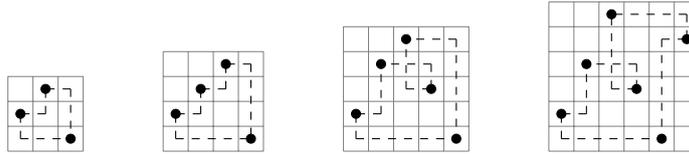


Figure 10: The cycles 231, 2341, 24531 and 246315.

Now consider the leftmost leaf. It again is positive, in relative position 1, so we get 2341. Then continue to the negative node to the right of the node just considered. The two potential children of the previous, positive, node corresponded to positions 1 and 2 of the cycle, so this node occupies position 3. As it is negative, we insert a 3 after position 3 of the cycle, obtaining 24531, depicted third in figure 10.

Last we consider the right, negative child of the root. It is in relative position 5, so we insert 5 after the last position of the cycle, obtaining our final cycle, 246315.

Note in this case all 3 of the elements of the permutation on the off diagonal, the elements 2,3 and 5 (in positions 1,4 and 6 respectively) correspond to leaf nodes. This is not the case for all representations of the tree, after certain rotations only some of these off-diagonal points still correspond to leaves.

From the description it isn't immediately clear this construction produces the same cycle regardless of the order in which nodes are processed, even before potential tree rotations are taken into account. We will show that the bijection is well-defined on a fixed tree diagram. First, to reduce the number of cases we must consider later, we demonstrate a relationship between the cycle constructed from a rooted-signed-binary tree and the one obtained from the "negative" of that tree.

Lemma 9. *Let T be a rooted-signed-binary tree and \bar{T} the rooted-signed-binary tree obtained from T by reversing the sign of each node. If σ is the cycle produced by processing the nodes of T in some fixed order, then processing the (oppositely signed) nodes of \bar{T} in the same order produces σ^{-1} the functional inverse of σ .*

Proof. The statement is certainly true when T consists of only the root node. As the root is unsigned, $T = \bar{T}$ and the associated cycle 21 is its own inverse.

Before continuing, in the cycle diagram of σ suppose that for indices i, j , with $i < j$, there is a consecutive vertical-horizontal pair of segments between diagonal points (i, i) and (j, j) . By definition, if the segments are above the diagonal then $\sigma(i) = j$; if they are below the diagonal then $\sigma(j) = i$. It follows that the cycle diagram obtained by reflecting the cycle diagram of σ across $y = x$ corresponds to the permutation σ^{-1} .

Now, suppose the statement holds for any rooted-signed-binary tree of size n . Let T be a tree of size $n + 1$, and v a leaf of T . Let σ be the permutation obtained from processing the nodes of $T - \{v\}$ in some order. By induction, $\bar{T} - \{v\}$ corresponds to σ^{-1} (processing the oppositely-signed nodes in the same order) and the corresponding cycle diagram is the reflection across the diagonal.

Exchanging the role of T and \overline{T} if necessary, assume v is positive and inserted in position i . Thinking of the cycle diagram of σ as determined by its n non-diagonal points $(i, \sigma(i))$, the effect of inserting v into $T - \{v\}$ is to move all points with height greater than i up one position, and shift all points in horizontal position at least i to the right by one. Additionally, a point at $(i, i + 1)$ is inserted. On the other hand, inserting a negative vertex in position i of $\overline{T} - \{v\}$ shifts points in the diagram of σ^{-1} that are in horizontal position greater than i to the right, and moves points with height at least i up one position. Additionally, a point at $(i + 1, i)$ is added.

We know that points which were unchanged, from the diagram corresponding to $T - \{v\}$ to the one for T , reflect across the diagonal to points of the diagram for $\overline{T} - \{v\}$. The observations above explain why those points which were shifted (or added) upon insertion of v will reflect across the diagonal to those which shift (or are added) when \bar{v} is inserted. Thus the permutations corresponding to T and \overline{T} (processing nodes in that fixed order) are functional inverses. \square

Proposition 10. *The cycle produced by applying the above construction to the nodes of a fixed representation of a rooted-signed-binary tree is the same, regardless of the order in which the nodes of the tree are processed.*

In order to prove Proposition 10 we will first prove the following.

Lemma 11. *The cycle produced by the above construction is not changed when the order of processing the last two leaves of the tree changes.*

Proof. Suppose T is a rooted-signed-binary tree with at least two leaves with $n + 2$ positions where a new leaf could be inserted. Fix two distinct leaves, v and w of T in relative positions i and j respectively of $T - \{v, w\}$ where $i < j$.

By processing the nodes of $T - \{v, w\}$ in some order, our construction assigns a cycle σ to $T - \{v, w\}$. Write $\sigma = s_1 s_2 \cdots s_n$ (in one line notation). We show v and w can be added to the tree in either order, and the corresponding changes to σ produce the same cycle.

First note we can assume by Lemma 9 that v is signed negatively (if it weren't we could swap T for \overline{T} and show that the inverse cycle can be constructed unambiguously regardless of the order in which the leaves are added). Recall the function ξ_m defined in (2). For any two integers $l < m$ it is easily verified that

$$\xi_l(\xi_m(k)) = \xi_{m+1}(\xi_l(k)) \tag{3}$$

holds for all integers k . We consider two cases, based on the sign of w .

Case 1: w is positive. Inserting the negatively-signed v first, in position i of $T - \{v, w\}$, we transform the associated permutation to

$$s_1 s_2 \cdots s_n \rightarrow \xi_i(s_1) \xi_i(s_2) \cdots \xi_i(s_i) [i] \xi_i(s_{i+1}) \cdots \xi_i(s_n). \tag{4}$$

Since w is to the right of v , w is now in relative position $j + 1$ of $T - \{w\}$. Also $\xi_i(s_j)$ is the $(j + 1)$ -element in (4), so inserting the positively signed w results in

$$\begin{aligned} & \cdots \xi_{j+2}(\xi_i(s_i)) \xi_{j+2}(i) \xi_{j+2}(\xi_i(s_{i+1})) \cdots \xi_{j+2}(\xi_i(s_{j-1})) [j+2] \xi_{j+2}(\xi_i(s_j)) \cdots \\ & = \cdots \xi_i(\xi_{j+1}(s_i)) [i] \xi_i(\xi_{j+1}(s_{i+1})) \cdots \xi_i(\xi_{j+1}(s_{j-1})) [j+2] \xi_i(\xi_{j+1}(s_j)) \cdots \end{aligned} \tag{5}$$

where the second line was obtained using (3) to interchange the two functions. On the other hand, if we first add w to $T - \{v, w\}$, the first permutation obtained is

$$s_1 s_2 \cdots s_n \rightarrow \xi_{j+1}(s_1) \xi_{j+1}(s_2) \cdots \xi_{j+1}(s_{j-1}) [j+1] \xi_{j+1}(s_j) \cdots \xi_{j+1}(s_n).$$

Inserting the negatively signed v into position i of $T - \{v\}$ is then seen to immediately yield (5) since $\xi_i(j+1) = j+2$.

Case 2: w is negative. As in case 1, if we first insert the negatively-signed node v we obtain (4). Again, w is to the right of v in T , so w is now in relative position $j+1$ of $T - \{w\}$. Inserting the negatively signed w in position $j+1$ results in

$$\begin{aligned} & \cdots \xi_{j+1}(\xi_i(s_i)) \xi_{j+1}(i) \xi_{j+1}(\xi_i(s_{i+1})) \cdots \xi_{j+1}(\xi_i(s_j)) [j+1] \xi_{j+1}(\xi_i(s_{j+1})) \cdots \\ & = \cdots \xi_i(\xi_j(s_i)) [i] \xi_i(\xi_j(s_{i+1})) \cdots \xi_i(\xi_j(s_j)) [j+1] \xi_i(\xi_j(s_{j+1})) \cdots \end{aligned} \quad (6)$$

interchanging the two functions using (3). If, instead, we first add w to $T - \{v, w\}$, the first, intermediate, permutation obtained is

$$s_1 s_2 \cdots s_n \rightarrow \xi_j(s_1) \xi_j(s_2) \cdots \xi_j(s_j) [j] \xi_j(s_{j+1}) \cdots \xi_j(s_n).$$

Now, inserting the negatively signed v into position i of $T - \{v\}$ again transforms this to (6) after noting that $\xi_i(j) = j+1$. \square

Proof of Proposition 10. Suppose for contradiction there exist trees for which the cycle construction described is ambiguous, and let T be a minimal-sized counterexample. Say that T has n nodes, and that processing these nodes in two orders, given by v_1, v_2, \dots, v_n and w_1, w_2, \dots, w_n , results in two different cycles. It must be that v_n and w_n are both leaves of T ; also $v_n \neq w_n$, since otherwise the cycle construction on $T - \{v_n\}$ is ambiguous.

Because T is a minimal counterexample, any order of processing the nodes of $T - \{v_n\}$ must result in the same cycle. And so we can assume $w_n = v_{n-1}$ (i.e. w_n is the last node inserted into $T - \{v_n\}$). Likewise, we may assume in the order w_1, w_2, \dots, w_n that $v_n = w_{n-1}$. Thus, each ordering first constructs a cycle for $T - \{v_n, w_n\}$ and by minimality the two cycles so constructed on $T - \{v_n, w_n\}$ must agree. Since, by assumption, processing T in the order v_1, v_2, \dots, v_n gives a different cycle than the order w_1, w_2, \dots, w_n , we have contradicted the statement of Lemma 11. \square

Lemma 12. *Suppose σ is a cycle of length at least 3 obtained from a rooted-signed-binary tree T . If $\sigma(i) - i = 1$, then T can be transformed by a series of rotations to a tree T' with a positively signed leaf in relative position i or, if $\sigma(i) - i = -1$, a negatively signed leaf in relative position $i - 1$.*

Proof. Note that if $\sigma(i) - i = -1$ then $\sigma^{-1}(i-1) - (i-1) = 1$. By Lemma 9 we can reduce to the case $\sigma(i) - i = 1$ by considering the inverse cycle with the oppositely signed tree. Hence, we consider only the case $\sigma(i) - i = 1$.

The cycles of length 3 are 231 and 312. Only the former has $\sigma(i) - i = 1$ (for $i = 1, 2$). In this case T has one positive leaf, which can be rotated into either position 1 or 2. Now, suppose the claim holds for all cycles of length $n - 1$. Let σ have length n , $\sigma(i) - i = 1$,

T a tree corresponding to σ with a leaf w in relative position k . If $k = i$ and w were positive we are done. If $k = i$ with w negative then $\sigma(i + 1) = i$. But $\sigma(i) = i + 1$ by assumption, and since σ is a cycle we would be left with $\sigma = 21$. This is not the case so we can assume $k \neq i$.

Let σ^* be the cycle corresponding to $T - \{w\}$. If w is negative and $k = i + 1$, then $\sigma(i + 2) = i + 1 = \sigma(i)$ which is impossible. Thus it is straightforward to check that if $k > i$ then $\sigma^*(i) = \sigma(i) = i + 1$ and, respectively, if $k < i$ that $\sigma^*(i - 1) = \sigma(i) - 1 = i$, regardless of the sign of w . Since σ^* has length $n - 1$, there exists a rotation of $T - \{w\}$, call it $(T - \{w\})'$, containing a positive leaf u in relative position i (respectively $i - 1$) of $(T - \{w\})'$. Note the children of u would be in relative positions i and $i + 1$ (respectively $i - 1$ and i).

If either $k < i - 1$ or $k > i + 1$ then u would still be a leaf after w is inserted into position k of $(T - \{w\})'$ and, in either case, u will be in position i . Since tree rotations do not affect the relative position of any leaf (so long as the rotation doesn't cause the node to no longer be a leaf) applying the same rotations to T required to transform $T - \{w\}$ into $(T - \{w\})'$ will result in a tree having a leaf in the desired position.

If $k = i - 1$ and w were negative, then $\sigma(i) = i - 1$ which isn't the case, so if $k = i - 1$ or $k = i + 1$ then w is positive, and a child of the leaf u in T' . In either case we can perform a tree rotation of w into u producing a positively signed vertex in position i as desired, and the same argument as above applies. \square

Proposition 13. *The construction of cycles from rooted-signed-binary trees is injective; trees that aren't related by tree rotations don't produce the same cycle.*

Proof. This is clear if $\sigma = 21$.

Let $n \geq 2$. Suppose that for any cycle of length n , and for any two rooted-signed-binary trees that produce this cycle, the two trees are related by tree rotations. Now, let σ be a cycle of length $n + 1$, and let T and T' be rooted-signed-binary trees that produce σ . We will show T and T' are related by tree rotations.

Choose a leaf v of T ; say that v is in relative position i . Note that T and T' are related by tree rotations if and only if \overline{T} and $\overline{T'}$ are related by tree rotations. Thus, by taking the negative trees \overline{T} and $\overline{T'}$ and renaming, if necessary, we may assume that v is a positive node.

As v is positive, the construction of σ is such that $\sigma(i) = i + 1$. Since T' produces σ as well, Lemma 12 implies that there is a tree that is related to T' by tree rotations, which has a positively signed leaf w in relative position i . And so, we may as well assume that T' has a positive leaf w in relative position i .

Let σ^* denote the cycle, with length n , produced from $T - \{v\}$. The procedure for how we obtain σ from σ^* , upon inserting v in position i of $T - \{v\}$, is reversible. Thus, since w is positive, in relative position i in T' , and T' produces σ , it must be that the tree $T' - \{w\}$ produces the cycle σ^* .

By assumption, there is a sequence of tree rotations taking $T - \{v\}$ to $T' - \{w\}$. Since v is not present, each of the rotations consists of a child node (that is *not* v) being rotated into the place of its parent. Therefore, the same sequence of rotations can be performed

on T and, in the result, the leaf v is still in the relative position i (recall Remark 5). That is, these tree rotations have transformed T into T' . □

So far we have not considered the allowed tree rotations of a rooted-signed-binary tree. In this section we show that two such trees which can be obtained from one another by allowed tree rotation moves correspond to the same cycle.

Theorem 14. *Any rooted-signed-binary-trees T and T' related by the rotations of Definition 2 produce the same cycle under the construction of Section 4.*

Proof. It suffices to show that any single tree rotation results in the same cycle, which we will prove by induction on the number of nodes in T . The result holds vacuously if T consists only of the unsigned root node. Suppose now the result holds for all trees smaller than T , and suppose T' is related to T by a tree rotation at some vertex v . By swapping T and T' if necessary, we suppose this rotation is a “clockwise” rotation, which rotates the left child of v into the place of v and creates a new right child of that vertex in T' .

Note, v and its left child must have the same sign. If this sign is negative we can replace T and T' with \bar{T} and \bar{T}' respectively so (by Lemma 9) we can suppose v and its left child are both positive (v could be the root, which is unsigned, but we still assume that its left child is positive, which doesn’t change the argument that follows.) We can draw the relevant portion of T and T' (showing only the vertex v and its descendants) as in Figure 7, where S_1 , S_2 and S_3 are the relevant subtrees for descendants of v and its left child.

Suppose w is a leaf appearing in one of these subtrees. Since w is not involved in the rotation it is a leaf of both T and T' . The trees $T - \{w\}$ and $T' - \{w\}$ are related by the same tree rotation at v and, since both are smaller than T , both correspond to the same cycle σ , by induction. Since w appears in the same relative position in both T and T' , adding w to either of $T - \{w\}$ or $T' - \{w\}$ would produce the same change to σ , thus T and T' correspond to the same cycle.

This leaves us with the case where all the subtrees S_1 , S_2 and S_3 are empty. Then the vertex v in T has a single positive left child w (with no children), while in T' the vertex w (now rotated into the place v occupied) only has a right child, say v' . Note that $T - \{w\}$ is the exact same tree as $T' - \{v'\}$. Let i be the relative position of v in $T - \{w\}$ and let σ be the cycle associated to $T - \{w\}$.

Since v is a positive leaf of $T - \{w\}$ we can write

$$\sigma = s_1 s_2 \cdots s_{i-1} [i + 1] s_{i+1} \cdots s_n$$

where $s_j \neq i + 1$. Now, the left and right children of v correspond to the positions i and $i + 1$ respectively. We check what occurs to the cycle σ when a positive node is inserted in either place. If a positive left child is inserted (creating T) then our rules dictate that an $i + 1$ is inserted in position i of σ , and all elements of sigma with value $i + 1$ or higher

are increased by 1. Thus this insertion results in

$$\begin{aligned} & \xi_{i+1}(s_1)\xi_{i+1}(s_2)\cdots\xi_{i+1}(s_{i-1}) [i+1] \xi_{i+1}(i+1)\xi_{i+1}(s_{i+1})\cdots\xi_{i+1}(s_n) \\ & = \xi_{i+1}(s_1)\xi_{i+1}(s_2)\cdots\xi_{i+1}(s_{i-1}) [i+1][i+2] \xi_{i+1}(s_{i+1})\cdots\xi_{i+1}(s_n). \end{aligned} \quad (7)$$

If a positive vertex is inserted as the right child of v (position $i+1$), creating T' , then our rules dictate that an $i+2$ is inserted in position $i+1$ of σ , and all elements of sigma with value $i+2$ or higher is increased by 1. Thus this insertion results in

$$\begin{aligned} & \xi_{i+2}(s_1)\xi_{i+2}(s_2)\cdots\xi_{i+2}(s_{i-1})\xi_{i+2}(i+1) [i+2] \xi_{i+2}(s_{i+1})\cdots\xi_{i+2}(s_n) \\ & = \xi_{i+2}(s_1)\xi_{i+2}(s_2)\cdots\xi_{i+2}(s_{i-1}) [i+1][i+2] \xi_{i+2}(s_{i+1})\cdots\xi_{i+2}(s_n). \end{aligned} \quad (8)$$

Since $\xi_{i+1}(s_j) = \xi_{i+2}(s_j)$ so long as $s_j \neq i+1$, we see that (7) and (8) are the same expression, and so T and T' correspond again to the same cycle. \square

5 Surjectivity

In this section we prove that the map from rooted-signed-binary trees to unknotted cycles is surjective. The proof relies on Bennequin's inequality [2], or more precisely a reformulation of it for Legendrian knots [12] – an important early result in modern contact geometry. We begin with some notation.

Definition 15. Let σ be a derangement. Define (i, j) to be a C -pair if either

$$i < j < \sigma(i) < \sigma(j) \quad \text{or} \quad i > j > \sigma(i) > \sigma(j).$$

Additionally, define i to be a UR -index if

$$\sigma^{-1}(i) < i \quad \text{and} \quad \sigma(i) < i.$$

Remark 16. Note that (i, j) is a C -pair if and only if two segments of the cycle diagram cross each other, one of the segments occurring between diagonal points (i, i) and $(\sigma(i), \sigma(i))$ and the other occurring between (j, j) and $(\sigma(j), \sigma(j))$. Such a pair appears as depicted in the left-most image of Figure 11 (or its reflection across the diagonal). All crossings in cycle diagrams have this form.

Also, i is a UR -index if and only if the cycle diagram has an upper right corner at (i, i) , as in the right-most image of Figure 11.

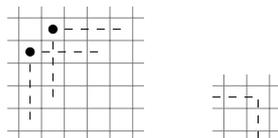


Figure 11: The cycle diagram at a C -pair (left) and UR -index (right).

Proposition 17 (Bennequin-Eliashberg Inequality). *Let D be a cycle diagram and let $K(D)$ be the knot associated to D . Define $C(D)$ to be the number of C -pairs of D and $UR(D)$ the number of UR -indices. Recalling that $g(K)$ is the genus of K , then*

$$C(D) - UR(D) \leq 2g(K(D)) - 1.$$

Proof. Let $\Lambda(D)$ be the Legendrian defined in Section 2. Cycle diagrams have all negative crossings, and so recall that each crossing of the front projection of $\Lambda(D)$ is positive, and so its writhe is $C(D)$. Points on D at a UR -index correspond to right cusps, and so

$$C(D) - UR(D) = \text{tb}(\Lambda(D)).$$

Recalling that the mirror of $K(D)$ is equivalent to $\Lambda(D)$ (simply as links), and that genus does not change under mirroring, the Bennequin-Eliashberg inequality [12] gives $\text{tb}(\Lambda(D)) \leq 2g(K(D)) - 1$. \square

Remark 18. In fact, for cycle diagrams D representing a knot K we have that $\text{tb}(\Lambda(D)) = 2g(K) - 1$, as the Seifert surface obtained directly from D has negative Euler characteristic agreeing with $\text{tb}(\Lambda(D))$. Yet, the inequality suffices for our purposes (see also [17]).

Proposition 19. *Suppose that σ is an n -cycle where $|\sigma(i) - i| \geq 2$ for all $1 \leq i \leq n$. Then the number of C -pairs of σ is at least as many as the number of UR -indices.*

Let D_σ be the cycle diagram of σ . Our proof uses the Seifert circles of D_σ , introduced in Section 2.2. The Seifert circles provide a convenient combinatorial device for our proof; knot theoretic constructions associated with them are not used. Recall the orientation convention for D_σ (found in Section 2). To each Seifert circle S associate a set $\mathcal{C}(S)$ of crossings of D_σ — the crossings which, when smoothed, created a portion of the Seifert circle.

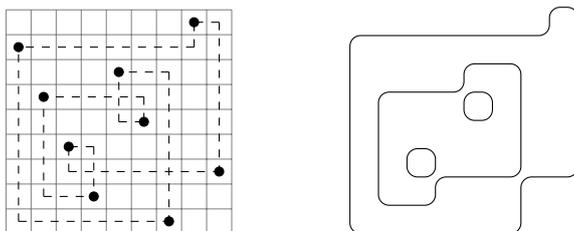


Figure 12: The cycle diagram of the unknotted cycle 864275193 and the corresponding Seifert circles obtained by smoothing the three crossings.

To prove the proposition, we will use few lemmas.

Lemma 20. *The number of Seifert circles of D_σ is equal to the number of UR -indices of D_σ .*

Proof. First, make a key observation: due to the orientation on D_σ , smoothing a crossing cannot *create* a local extremum of the function $x + y$ on a Seifert circle. Therefore, any such local extremum must correspond to a point on D_σ itself. So it is a diagonal point, at a UR -index if it is a local maximum.

As a Seifert circle is a closed planar curve, it has at least one local maximum of $x + y$, and so passes through at least one diagonal point that is at a UR -index.

The Seifert circles of D_σ have an induced orientation that agrees with the orientation on D_σ away from the crossings (oriented clockwise in the case of cycle diagrams). If S is a Seifert circle of D_σ , then at each point where S passes through the diagonal, it goes from above the diagonal to below it. Hence, S cannot pass through two UR -indices since, were it to do so, it would spiral either inward or outward and be unable to form a simple closed curve. Thus, every Seifert circle passes through exactly one UR -index. \square

Remark 21. Similarly, there is exactly one lower-left corner on the diagonal (an LL -index, say) contained in any Seifert circle of a cycle diagram.

Now, place a partial ordering on the set of Seifert circles. If S and S' are Seifert circles, then we say that $S' \prec S$ when S' is contained in the bounded planar region that is enclosed by S .

Lemma 22. *There is a unique maximal Seifert circle of D_σ , larger than every other Seifert circle.*

Proof. Let S be a maximal Seifert circle (in some maximal chain). If there exist Seifert circles in the unbounded region of S , then at least one such circle S' must share a crossing with S . Otherwise it would not be a knot (σ would not be a cycle).

Now, suppose that the UR -index for S is less than the UR -index for S' . Since $S \not\prec S'$, the UR -index of S is also less than the unique LL -index of S' . However, this makes it impossible for there to be a crossing in $\mathcal{C}(S) \cap \mathcal{C}(S')$. Likewise, if we suppose that the UR -index of S is greater than that of S' , then the LL -index of S must also be larger, making a shared crossing impossible.

Hence, every other Seifert circle is in the bounded region of S . \square

Lemma 23. *If D_σ has a crossing, then any Seifert circle S has at least 1 associated crossing. Also, under the hypothesis of Proposition 19, if $|\mathcal{C}(S)| = 1$, then S is the maximal Seifert circle.*

Proof. The first statement is clear, since otherwise S is a closed curve already in the (unsmoothed) cycle diagram and, as the diagram has a crossing elsewhere, there would need be multiple components (a link, rather than a knot).

Suppose $\mathcal{C}(S)$ contains only one crossing. It must be that either S is a minimal element in the partial order, or it is maximal. If this weren't so, then there would be a Seifert circle in both the bounded and unbounded regions of S . But the fact that the knot has a single component then necessitates another crossing in $\mathcal{C}(S)$.

Now, suppose S is minimal. Let ℓ be the LL -index of S and let r be the UR -index (which is unique by Lemma 20). Since S is minimal, S must intersect every diagonal

point between ℓ and r . We can divide S into segments, each having one vertical part and one horizontal part. Each segment will start and end on the diagonal: for some i and j the segment has a vertical segment from (i, i) to (i, j) followed by a horizontal segment to (j, j) . Call such a segment a *step* of S , with length $|i - j|$.

If any step of length 1 occurs in S , from (i, i) to $(i + 1, i + 1)$ for some i , then there is a smoothed crossing that was located at $(i, i + 1)$; likewise, if there is a step of length 1 from $(i + 1, i + 1)$ to (i, i) , then prior to smoothing out crossings there was a crossing located at $(i + 1, i)$ – a step of length 1 could not have existed in D_σ by the hypothesis of Lemma 19. So if we show that S has at least two steps of length 1, then we have a contradiction. Note that this is automatic if $r - \ell = 1$ (with one step above the diagonal and one below), so we assume that $r - \ell > 1$.

For any i with $\ell < i < r$, the point (i, i) is part of two steps of S , either both above the diagonal or both below. Now consider the point $(\ell + 1, \ell + 1)$. As it is part of two steps (and is not an LL -index), it must be connected to (ℓ, ℓ) , a step of length 1. A similar argument holds for $r - 1$ and r , giving another step of length 1. These steps are not the same as $r - \ell > 1$. By the above argument, $|\mathcal{C}(S)| > 1$, a contradiction.

Therefore, S cannot be minimal in the partial order when $|\mathcal{C}(S)| = 1$, and S is the (unique) maximal Seifert circle. \square

Proof of Proposition 19. By Lemma 20, we must show that the the number of Seifert circles of D_σ is at most the number of crossings. Let s denote the number of such Seifert circles. Since each crossing is associated to two Seifert circles, and by Lemma 23 every Seifert circle S has $|\mathcal{C}(S)| \geq 2$ (except possibly the maximal Seifert circle), the total number of crossings must be at least $\lceil \frac{2s-1}{2} \rceil = s$. \square

If K is an unknot, then it bounds a disk and so the Seifert genus is $g(K) = 0$. Hence we obtain the following.

Proposition 24. *If the link associated to the cycle diagram for σ is the (single component) unknot, then there is some i such that $|i - \sigma(i)| = 1$.*

Proof. Let D_σ be the cycle diagram. Since $g(K) = 0$, Proposition 17 implies that $C(D_\sigma) \leq UR(D_\sigma) - 1$. By Lemma 19 we must have some i such that $|i - \sigma(i)| < 2$. The fact that σ is a derangement thus implies the conclusion. \square

Proposition 25. *Every unknotted cycle is obtained from a rooted-signed-binary tree through the construction of Section 4.*

Proof. The proof will be by induction on the length of the cycle. A cycle of length 2 it must be $\sigma = 21$ which corresponds to the tree containing only the root.

Let D be the cycle diagram of an unknotted cycle σ of length n . By Proposition 24, there is an i , with $1 \leq i \leq n$ so that $|i - \sigma(i)| = 1$. If $i - \sigma(i) = 1$ then, near the i^{th} diagonal point, D must be like the depiction in one of the rows in the third column of Table 1. Let D_0 be the cycle diagram obtained by collapsing the i and $i - 1$ -st rows and columns of D , replacing them with what appears in the same row, but first column of 1. (Thus, D_0 is a cycle diagram with one fewer row and column than D .) If $i - \sigma(i) = -1$

instead, then D must be like the depiction in one of the rows of the second column of Table 1. In similar manner to the previous case, define D_0 by using the same row, but in the first column.

As it has fewer diagonal nodes, we may assume that D_0 is obtained from a rooted-signed-binary tree T_0 . Now, define a tree T by adding a node to T_0 : if $i - \sigma(i) = 1$ then add a negative node to T_0 at position $i - 1$; if $i - \sigma(i) = -1$ then add a positive node to T_0 at position i . Then clearly T is assigned cycle diagram D by the construction. \square

6 Further Results

Extending these ideas, we can count all unlinked permutations (derangements). By an unlink, we mean the knot corresponding to each cycle of the permutation is an unknot in the grid diagram, and furthermore each of these unknots are not “linked” with one another. We obtain a bivariate generating function that keeps track of both the size of the permutation and the number of knots in the link.

Theorem 26. *Let \mathcal{U} be the set of unlinked derangements and denote by $\text{cyc}(\sigma)$ the number of cycles in σ (equivalently knots in the link associated to σ). Define the bivariate generating function*

$$F(u, x) = 1 + \sum_{\sigma \in \mathcal{U}} u^{\text{cyc}(\sigma)} x^{|\sigma|} \quad (9)$$

to count unlinked permutations by length and number of components. Then $F(u, x)$ satisfies the recurrence

$$2 + (ux - 2)F(u, x) - ux^2F(u, x)^2 - uxF(u, x)\sqrt{1 - 6xF(u, x) + x^2F(u, x)^2} = 0 \quad (10)$$

or, equivalently,

$$1 + (ux - 2)F(u, x) + (1 - ux - ux^2)F(u, x)^2 + (ux^2 + u^2x^3)F(u, x)^3 = 0.$$

Setting $u=1$ we recover the generating function $f(x)$ for the sequence of all unlinked permutations, and find that $f(x)$ satisfies the recurrence

$$1 + (-2 + x)f(x) + (1 - x - x^2)f(x)^2 + (x^2 + x^3)f(x)^3 = 0$$

corresponding to the sequence 1, 2, 8, 32, 143, 674, 3316, 16832, 87538 This is listed, along with the sequences counting unlinked permutations with k disjoint cycles for each $k < 5$, in Table 2. Note that the number of unlinked permutations of $2n$ with n components is the n -th Catalan number, which is easy to prove by relating the nested links to Dyck paths.

To prove Theorem 26 we need the following result.

Lemma 27. *Let σ be an unlinked permutation, with cycle decomposition $\sigma = \sigma_1\sigma_2 \cdots \sigma_k$. Then each of the component cycles σ_i are unknotted cycles, and the cycle diagrams of any two cycles are noncrossing.*

k	# k -component unlinked permutations	Generating Function
1	0, 1, 2, 6, 22, 90, 394, 1806, 8558, 41586 ...	$\frac{1}{2}(x - x^2 - x\sqrt{x^2 - 6x + 1})$
2	0, 0, 0, 2, 10, 48, 238, 1216, 6354, 33760 ...	$\frac{1}{2}x^2 - 3x^3 + x^4 - \frac{x^2 - 9x^3 + 12x^4 - 2x^5}{2\sqrt{x^2 - 6x + 1}}$
3	0, 0, 0, 0, 0, 5, 42, 280, 1752, 10710
4	0, 0, 0, 0, 0, 0, 0, 14, 168, 1440
all	0, 1, 2, 8, 32, 143, 674, 3316, 16832, 87538 ...	$f(x)$

Table 2: The number of unlinked permutations by the number of components in the link.

Proof. We show that the cycle diagrams of two different component cycles of an unlinked permutation cannot cross each other by considering the linking number of the corresponding unknots. Let K and K' be two different components of a link L , and consider a diagram in the plane of L . Define c_+ (resp. c_-) to be the number of positive (resp. negative) crossings in the diagram, where we only count crossings that involve one strand from K and one strand from K' . The *linking number* of K and K' equals $1/2(c_+ - c_-)$.

For any planar diagram that represents a link equivalent to L there will be two components corresponding to K and K' . It is well-known (see e.g. [19, pg. 11]) that the linking number (in that diagram) of those two components must be equal to the previously computed number, $1/2(c_+ - c_-)$. That is, the linking number of K and K' is invariant under change of diagram. (In fact, one can verify this by checking that the linking number will be unchanged by Reidemeister moves.)

Since there is a diagram of the $cyc(\sigma)$ component unlink which has no crossings at all, the linking number of any two components in any diagram must be zero. However, all crossings are negative in cycle diagrams. If two component cycle diagrams cross at any point, then the linking number of those components cannot be zero, and so the diagram cannot be that of an unlinked permutation. \square

Define the support of a length n permutation σ , $\text{sup}(\sigma) = \{1 \leq i \leq n \mid \sigma(i) \neq i\}$.

Lemma 28. *Suppose σ is a derangement of n whose cycle diagram contains no crossings between different components of its cycle diagram. Let σ' denote the cycle of σ containing 1 and let $1 = \sigma'_1 < \sigma'_2 < \dots < \sigma'_j$ be the elements of $\text{sup}(\sigma')$. Then there exist permutations $\tau_1, \tau_2, \dots, \tau_j$ with*

$$\sigma = \sigma' \circ \tau_1 \circ \tau_2 \circ \dots \circ \tau_j \tag{11}$$

such that the support of τ_i is precisely the integers between σ'_i and σ'_{i+1} ,

$$\text{sup}(\tau_i) = \{k \mid \sigma'_i < k < \sigma'_{i+1}\}, \text{ for } i < j \text{ and } \text{sup}(\tau_j) = \{k \mid \sigma'_j < k \leq n\}. \tag{12}$$

Example 29. Take $\sigma = 732541698$ (one line notation) as depicted in Figure 13. Written in cycle notation $\sigma = (176)(23)(45)(89)$. Then (in cycle notation) $\sigma' = (176)$, $\sigma'_1 = 1, \sigma'_2 = 6, \sigma'_3 = 7$ and $\tau_1 = (23)(45)$, τ_2 is the identity, and $\tau_3 = (89)$.

Proof of Lemma 28. For each $1 \leq i \leq j$, define τ_i to be the permutation consisting of all cycles from σ that contain in their support at least one element strictly between σ'_i and

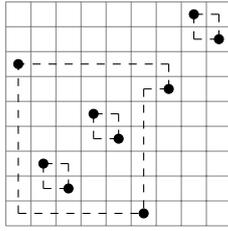


Figure 13: $\sigma = 732541698$.

σ'_{i+1} (or strictly greater than σ'_j , in the case of τ_j). It suffices to show that every element in $\text{sup}(\tau_i)$ is contained in the desired range, between σ'_i and σ'_{i+1} (or greater than σ'_j , when $i = j$).

To this end, suppose to the contrary there were a $k < \sigma'_{i+1}$ with $\tau_i(k) > \sigma'_{i+1}$. (The case where $\tau_i(k) < \sigma'_i$ proceeds similarly.) While τ_i need not consist of a single cycle, there must exist some $k' \in \text{sup}(\tau_i)$ so that $k' > \sigma'_{i+1}$, and yet $1 < \tau_i(k') < \sigma'_{i+1}$ (the first iteration of τ_i , $n \geq 1$, so that $k' = \tau_i^n(k) > \sigma'_{i+1}$ but $\tau_i^{n+1}(k) < \sigma'_{i+1}$).

Recall that if $k < \sigma(k)$, the corresponding vertical and horizontal segments of the cycle diagram for σ must be *above* the diagonal, while if $k > \sigma(k)$, the corresponding segments must be *below* the diagonal. Thus, the elements k, k' determine a cycle component whose corresponding diagram surrounds the diagonal point $(\sigma'_{i+1}, \sigma'_{i+1})$. That component's diagram is a closed curve in the plane, as depicted in Figure 14. Since elements in the support of τ_i must all be greater than 1, the point $(1, 1)$ lies in the unbounded region for this closed curve, and the diagram for σ' must reach the bounded region to arrive at $(\sigma'_{i+1}, \sigma'_{i+1})$. There must then be a crossing between these components, contradicting the assumption there is no such crossing. \square

Finally, we count unlinked permutations.

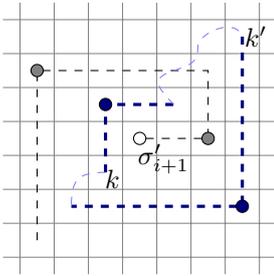


Figure 14: The cycle diagram lines for τ_i .

Proof of Theorem 26. In the definition of the bivariate generating function $F(u, x)$ the initial term 1 accounts for an empty derangement, which will be useful in the recursion. For the remainder of the proof we assume we are counting nontrivial derangements, which necessarily have length at least 2.

We now construct all such unlinked derangements as follows. We select first the cycle containing the element 1. Once we fix the number of elements to be in this cycle, k , then by Theorem 1 the number of possible choices for a cycle supported on k elements is S_{k-1} . (Note that the knot type is unchanged when the numbers in the support of the cycle are changed, so long as their relative order is preserved, as shifting the values, without changing their order has the effect only of changing the lengths of the lines in the corresponding cycle diagram.)

Having chosen the number of elements permuted by the cycle containing 1, σ' as well as the relative cycle type on those elements, we can apply Lemma 28, which tells us that in between each of the elements of σ' we can insert any unlink (including potentially the empty unlink) which is counted by our original generating function $F(u, x)$. Since we can insert such an unlink between any of the k elements of σ' , as well as after the last element, there are a total of k places where such an unlink can be inserted. Adding together these possibilities, we have

$$\begin{aligned} F(u, x) &= 1 + \sum_{k=2}^{\infty} S_{k-1} u x^k F(u, x)^k = 1 + u x F(u, x) \sum_{k=1}^{\infty} S_k (x F(u, x))^k \\ &= 1 + u x F(u, x) S(x F(u, x)). \end{aligned}$$

Using that $S(x) = \frac{1}{2} (1 - x - \sqrt{1 - 6x + x^2})$ now gives

$$F(u, x) = 1 + \frac{u x}{2} F(u, x) (1 - x F(u, x) - \sqrt{1 - 6x F(u, x) + (x F(u, x))^2})$$

which simplifies to (10). □

Finally, one could instead choose to consider all permutations, rather than just derangements, with the convention that any fixed points in the permutation correspond to infinitesimal unknot components of the unlink. In this case, we obtain the following modification of Theorem 26 by the same argument.

Theorem 30. *Let \mathcal{V} be the set permutations whose cycle diagram corresponds to an unlink (treating fixed points as their own component of an unknot) and define the bivariate generating function*

$$G(u, x) = 1 + \sum_{\sigma \in \mathcal{V}} u^{cyc(\sigma)} x^{|\sigma|}. \quad (13)$$

Then $G(u, x)$ satisfies the recurrence

$$2 + (3ux - 2)G(u, x) - ux^2 G(u, x)^2 - ux G(u, x) \sqrt{1 - 6xG(u, x) + x^2 G(u, x)^2} = 0, \quad (14)$$

or equivalently

$$ux^2 G(u, x)^3 + (2u^2 x^2 - ux^2 - 3ux + 1)G(u, x)^2 + (3ux - 2)G(u, x) + 1 = 0.$$

Setting $u = 1$ the sequence counting such permutations with any number of components begins

$$1, 2, 6, 23, 103, 511, 2719, 15205, 88197, \dots$$

which appears to match entry [A301897](#) in the OEIS. This sequence counts permutations with the following property. Given a permutation σ , let $\text{inv}(\sigma)$ be the number of inversions,

$\text{cyc}(\sigma)$ the number of cycles, and $\text{td}(\sigma)$ the total displacement, defined by Diaconis and Graham [10] to be

$$\text{td}(\sigma) = \sum_{i=1}^{|\sigma|} |\sigma(i) - i|$$

(see also [22]). Diaconis and Graham prove that

$$\text{inv}(\sigma) + (|\sigma| - \text{cyc}(\sigma)) \leq \text{td}(\sigma). \quad (15)$$

The OEIS sequence above counts those permutations for which the Diaconis-Graham inequality is an equality. Recently, Alexander Woo has shown, using the results of this paper, that this set of permutations is the same as permutations that give an unlink [26]. Also, since this paper was first published on the arXiv, Berman and Tenner [3] showed that the “fundamental bijection” sending cycle notation to one-line notation restricts to a bijection between cycles for which (15) is an equality and the set of separable permutations.

References

- [1] S. Baader, L. Lewark, and L. Liechti, *Checkerboard graph monodromies*, Enseign. Math. **64** (2018), no. 1, 65–88.
- [2] D. Bennequin, *Entrelacement et équations de Pfaff*, Asterisque **107–108** (1983), 83–161.
- [3] Y. Berman and B. E. Tenner, *Pattern-functions, statistics, and shallow permutations*, 2021. [arXiv:2110.11146](https://arxiv.org/abs/2110.11146).
- [4] J. S. Birman and T. E. Brendle, *Chapter 2 - Braids: A survey*, Handbook of knot theory, 2005, pp. 19–103.
- [5] P. Bose, J. F. Buss, and A. Lubiw, *Pattern matching for permutations*, Inform. Process. Lett. **65** (1998), no. 5, 277–283.
- [6] M. Bouvel and D. Rossin, *The longest common pattern problem for two permutations*, Pure Math. Appl. **17** (2006), no. 1-2, 55–69.
- [7] P. R. Cromwell, *Homogeneous links*, J. London Math. Soc. (2) **39** (1989), no. 3, 535–552.
- [8] P. R. Cromwell, *Embedding knots and links in an open book. I. Basic properties*, Topology Appl. **64** (1995), no. 1, 37–58.
- [9] R. H. Crowell and R. H. Fox, *Introduction to knot theory*, Graduate Texts in Mathematics, No. 57, Springer-Verlag, New York-Heidelberg, 1977. Reprint of the 1963 original.
- [10] P. Diaconis and R. L. Graham, *Spearman’s footrule as a measure of disarray*, J. Roy. Statist. Soc. Ser. B **39** (1977), no. 2, 262–268.

- [11] I. A. Dynnikov, *Arc-presentations of links: monotonic simplification*, Fund. Math. **190** (2006), 29–76.
- [12] Y. Eliashberg, *Contact 3-manifolds twenty years since J. Martinet’s work*, Annales de l’institut Fourier **42** (1992), no. 1-2, 165–192 (eng).
- [13] S. Elizalde, *The X-class and almost-increasing permutations*, Ann. Comb. **15** (2011), no. 1, 51–68.
- [14] J. B. Etnyre, *Legendrian and transversal knots*, Handbook of knot theory, 2005, pp. 105–185.
- [15] C. Even-Zohar, *Models of random knots*, J. Appl. and Comput. Topology **1** (2017), 263–296.
- [16] F. Frankl and L. Pontrjagin, *Ein Knotensatz mit Anwendung auf die Dimensionstheorie*, Math. Ann. **102** (1930), no. 1, 785–789.
- [17] K. Hayden and S. J.M., *Positive knots and Lagrangian fillability*, Proc. Amer. Math. Soc. **143** (2015), 1813–1821.
- [18] D. E. Knuth, *The art of computer programming. Vol. 4A. Combinatorial algorithms. Part 1*, Addison-Wesley, Upper Saddle River, NJ, 2011.
- [19] W. B. R. Lickorish, *An introduction to knot theory*, Graduate Texts in Mathematics, vol. 175, Springer-Verlag, New York, 1997.
- [20] L. Ng and D. Thurston, *Grid diagrams, braids, and contact geometry*, Proceedings of Gökova Geometry-Topology Conference 2008, 2009, pp. 120–136.
- [21] O. E. of Integer Sequences, *Sequence A006318* (N. J. A. Sloane, ed.)
- [22] T. K. Petersen and B. E. Tenner, *The depth of a permutation*, J. Comb. **6** (2015), no. 1-2, 145–178.
- [23] H. Seifert, *Über das Geschlecht von Knoten*, Math. Ann. **110** (1935), no. 1, 571–592.
- [24] L. Shapiro and A. B. Stephens, *Bootstrap percolation, the Schröder numbers, and the N-kings problem*, SIAM J. Discrete Math. **4** (1991), no. 2, 275–280.
- [25] S. L. Witte, *Link nomenclature, random grid diagrams, and Markov Chain methods in knot theory*, 2020. Ph.D. Dissertation – University of California-Davis.
- [26] A. Woo, *The shallow permutations are the unlinked permutations*, 2022. [arXiv:2201.12949](https://arxiv.org/abs/2201.12949).