

REPRESENTING PERMUTATIONS WITH FEW MOVES

SERGEY BEREG, ALEXANDER E. HOLROYD,
LEV NACHMANSON, AND SERGEY PUPYREV

ABSTRACT. Consider a finite sequence of permutations of the elements $1, \dots, n$, with the property that each element changes its position by at most 1 from any permutation to the next. We call such a sequence a *tangle*, and we define a *move* of element i to be a maximal subsequence of at least two consecutive permutations during which its positions form an arithmetic progression of common difference $+1$ or -1 . We prove that for any initial and final permutations, there is a tangle connecting them in which each element makes at most 5 moves, and another in which the total number of moves is at most $4n$. On the other hand, there exist permutations that require at least 3 moves for some element, and at least $2n - 2$ moves in total. If we further require that every pair of elements exchange positions at most once, then any two permutations can be connected by a tangle with at most $O(\log n)$ moves per element, but we do not know whether this can be reduced to $O(1)$ per element, or to $O(n)$ in total. A key tool is the introduction of certain restricted classes of tangle that perform pattern-avoiding permutations.

1. INTRODUCTION

Let S_n be the symmetric group of permutations $\pi = [\pi(1), \dots, \pi(n)]$ on $\{1, \dots, n\}$, with composition defined via $(\pi \cdot \rho)(i) = \pi(\rho(i))$. It is natural to represent a permutation π as a composition of simpler permutations. Define the **swap** $s(i)$ to be the permutation $[1, \dots, i - 1, i + 1, i, \dots, n]$ that interchanges i and $i + 1$. We call two permutations π and ρ **adjacent** if they are related by a collection of non-overlapping swaps, i.e. if $\rho = \pi \cdot s(p_1) \cdots s(p_k)$ where $|p_i - p_j| \geq 2$ for $i \neq j$. Equivalently, π and ρ are adjacent if $|\pi^{-1}(i) - \rho^{-1}(i)| \leq 1$ for every i . A **tangle** is a finite sequence of permutations in which each consecutive pair is adjacent. If a tangle T starts with the identity permutation $\text{id} = [1, \dots, n]$ and ends with π , we say that T **performs** π .

It is straightforward to see that for any permutation π there is some tangle that performs π . Our goal is to find tangles with simple and elegant

Date: 11 August 2015.

2010 Mathematics Subject Classification. 05D99; 05A05; 68R10.

Key words and phrases. permutation, permutation diagram, reduced word, graph drawing, permutation pattern.

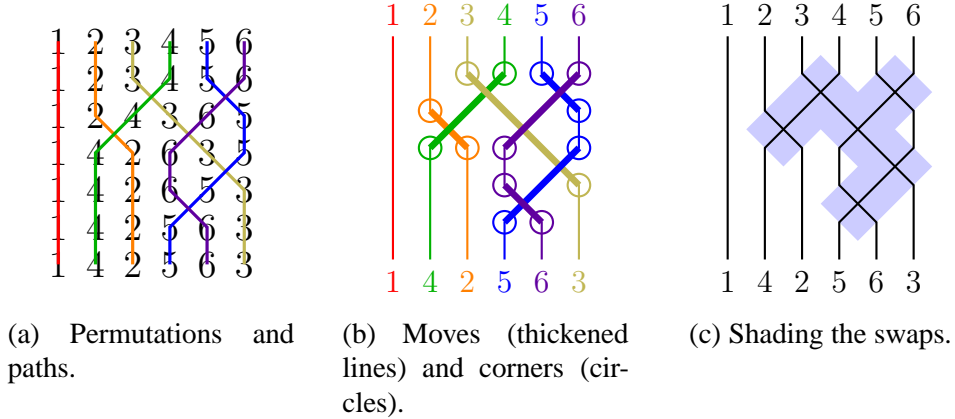


FIGURE 1. A tangle performing the permutation $\pi = [1, 4, 2, 5, 6, 3]$, with 7 moves.

structure. We may visualize a tangle as follows. Consider the sequence of permutations written in one-line notation $\pi = [\pi(1), \dots, \pi(n)]$ in a column from top to bottom as in Figure 1(a), with equal horizontal and vertical spacings between symbols. Then, for each $i = 1, \dots, n$, draw a polygonal path connecting all occurrences of the number i , from top to bottom, as in the figure. The path corresponding to element i is called **path** i . Each line segment of a path is either vertical or at an angle of $\pm 45^\circ$ to the vertical. We call a maximal non-vertical line segment of a path a **move**. Thus, a move corresponds to a maximal sequence of swaps $s(p_i)$ that occur between the adjacent elements in some interval of permutations of the tangle, and with their locations p_i forming an arithmetic progression with common difference ± 1 . See Figure 1(b). It is convenient to illustrate the structure by shading the area occupied by swaps, as in Figure 1(c). Our focus is on minimizing moves among tangles that perform a given permutation.

Our first main result is that any permutation can be performed by a tangle with a bounded number of moves per path (and therefore $O(n)$ moves in total as $n \rightarrow \infty$). In contrast, various natural greedy algorithms for constructing a tangle (including one proposed in [18]) require $\Omega(n^2)$ moves in total in the worst case. (See Figure 3 for examples.)

Theorem 1. *For any permutation $\pi \in S_n$, there is a tangle performing π that has at most 5 moves in each path.*

Shifting our attention to *total* moves, we can reduce the constant from 5 to 4.

Theorem 2. *For any permutation $\pi \in S_n$, there is a tangle performing π that has at most $4n$ moves in total.*

On the other hand, for all sufficiently large n there are permutations that require at least 3 moves in some path, and permutations that require at least $2n - 2$ moves in total. (The latter is easily seen to hold for the reverse permutation $[n, n - 1, \dots, 1]$, while the former apparently requires a quite involved argument – see Proposition 17). It is an open problem to close the gap between the bounds 3 and 5 for moves per path, and between $2n - 2$ and $4n$ for total moves.

Figures 2(a) and 2(b) give examples of the constructions behind Theorems 1 and 2. The tangles will be constructed by combining various “gadgets” – smaller tangles that are capable of performing permutations in certain restricted classes. Specifically, we will consider gadgets that perform (and are in bijective correspondence with) Grassmannian, 321-avoiding, 213-avoiding, and 132-avoiding permutations.

Despite the relatively small numbers of moves, the tangles illustrated in Figures 2(a) and 2(b) arguably have some undesirable features, which we discuss next. Firstly, they have many “holes” – small internal regions containing no swaps, shown unshaded in the figures. Secondly, a given pair of paths may cross multiple times. We will show that some version of the first issue is unavoidable if the number of moves is to be linear in n . On the other hand, we do not know whether the second issue can be avoided.

Rather than holes, it will be convenient to work with a slightly different notion, to be defined next. First we observe that counting moves is essentially equivalent to counting corners (see also [4]). A **corner** is a vertex of a path, at which its direction changes between any two of the three possible directions. Assume that a tangle has its initial and final permutations repeated at least once, so that each path starts and ends with a vertical segment. In addition, count “double corners” (at which a path changes from one non-vertical direction to the other) with multiplicity 2. With these conventions, the number of corners in a path equals twice the number of moves.

In our geometric interpretation of a tangle, we think of the swaps as located at the elements of the integer lattice \mathbb{Z}^2 . Therefore, the elements of the permutations, and thus also the corners, are located at elements of the shifted lattice $(\mathbb{Z} + \frac{1}{2})^2$. Specifically, take the i th element $\pi_t(i)$ of the t th permutation π_t in the tangle to be located at the point $(i - \frac{1}{2}, t - \frac{1}{2})$, where the first coordinate increases from left to right, and the second coordinate increases from top to bottom.

Given a tangle T , consider the graph whose vertices are the corners of T , and with an edge between two corners if their locations are within ℓ^∞ -distance 1. We call the connected components of this graph **clusters**. (See Figure 16.) The idea is that clusters generalize the notion of holes discussed above. Our next result implies that, as $n \rightarrow \infty$, for some (in fact, almost all)




FIGURE 2. Examples of the tangles corresponding to the main results. Shading is added to illustrate the structure.

permutations, if a tangle has only $O(n)$ corners (equivalently, $O(n)$ moves) then it must have at least $\Omega(n)$ clusters. Indeed, $o(n)$ clusters necessitates $\Omega(n \log n)$ corners. The proof will use a counting argument.

Theorem 3. *Let $\theta \in (0, \frac{1}{2})$ and suppose $n > \theta^{-8/\theta}$. For at least a proportion $1 - e^{-n}$ of the permutations $\pi \in S_n$, any tangle performing π has either at least $(\frac{1}{2} - \theta)n$ clusters or at least $\frac{1}{6}\theta n \log n$ corners.*

We now turn to the second issue raised above. We call a tangle **simple** if each pair of paths has at most one crossing. It is again easy to see that every permutation admits a simple tangle. In a simple tangle performing a permutation π , paths $\pi(i)$ and $\pi(j)$ cross each other if and only if $(\pi(i), \pi(j))$ is an **inversion** of π , i.e. $i < j$ and $\pi(i) > \pi(j)$.

The article [4] by the current authors characterizes a class of permutations for which there exist *simple* tangles that have the minimum moves among *all* tangles. However, there exist permutations that require strictly more moves for a simple tangle than for a general tangle. Again, see [4] for details.

In contrast with the case of general tangles discussed earlier, our upper and lower bounds for numbers of moves in simple tangles are rather far apart: $O(n \log n)$ and $\Omega(n)$ respectively as $n \rightarrow \infty$. Closing this gap is our principal open problem.

Proposition 4. *For any permutation $\pi \in S_n$, there is a simple tangle performing π that has at most $\lceil \log_2 n \rceil$ moves in each path.*

Proposition 5. *For every $n \geq 1$, there is a permutation $\pi \in S_n$ such that any simple tangle that performs it has at least $3n - c\sqrt{n}$ moves, where $c > 0$ is an absolute constant.*

While our focus is on moves, one can attempt to optimize other aspects of a tangle. For instance, we may define the **depth** of a tangle to be the length of the sequence of permutations comprising it (including the final permutation but not the initial one, say). It is not difficult to check that any $\pi \in S_n$ can be performed by some tangle of depth at most $n - 1$ for even n and at most n for odd n (and these bounds are optimal; they are attained by the reverse permutation). Our constructions for Theorems 1 and 2 and Proposition 4 perform reasonably well in this regard, having depths at most $3n$, $7n/4$ and $3n/2$ respectively.

Background. Further material on tangles and moves appears in a companion paper [4] by the current authors. The main result of [4] is a surprisingly complex characterization of the set of permutations that can be performed by a simple tangle in which each path has at most one move in each direction, together with a polynomial-time algorithm for recognizing such a

permutation and constructing the tangle. (In particular, this set turns out to include every permutation in S_6 , but no permutation containing the pattern 7324651.) Tangles and related objects have been studied in several settings by other authors, although the problem of minimizing moves (or corners) does not appear to have been considered prior to [4].

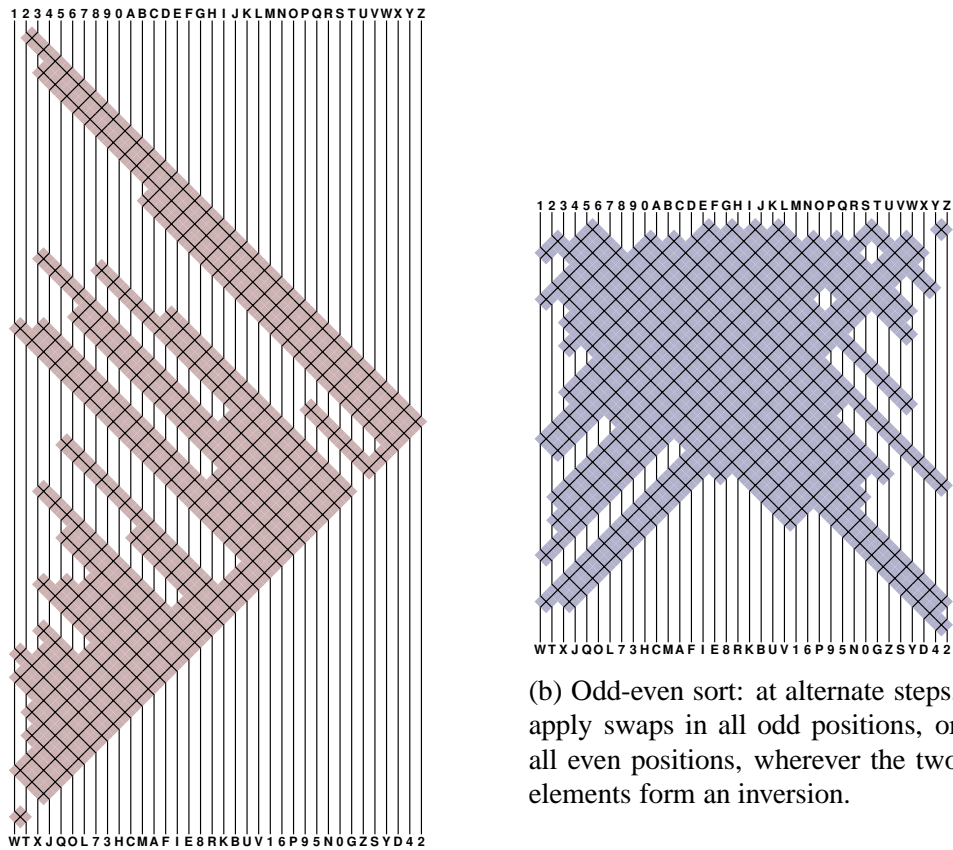
Wang in [18] considered essentially the same notion in the context of VLSI design for integrated circuits. However, the research in [18] targets, in our terminology, the depth of a tangle, and the total length of the paths. The algorithm suggested by Wang produces tangles with $O(n^2)$ moves for some permutations.

In algebraic combinatorics, Schubert polynomials can be encoded as sums over diagrams called RC-graphs or pipe dreams [5, 9], which may be interpreted as tangles of a certain type. Specifically, an RC-diagram corresponds via a 45° rotation to a simple tangle whose swaps are restricted to odd locations in a triangular region (the same region as our “reflector gadget” in Section 2.3). Reduced words for permutations are extensively studied; see e.g. [3, 10, 14, 21]. In our terminology, a reduced word is a simple tangle with only one swap between consecutive permutations.

Decomposition of permutations into nearest-neighbour transpositions was considered in the context of permuting machines and pattern-restricted classes of permutations [2]. In our terminology, Albert et. al. [2] proved that it is possible to check in polynomial time whether for a given permutation there exists a tangle of depth k , for a given k . Tangles and the associated visualizations also appear in sorting networks [1, 14], in arrangements of pseudolines [8], and in the context of change ringing (English-style church bell ringing) [20]. In the terminology of change ringing, a tangle with minimum corners is a “link method with minimum changes of direction”; each permutation represents an order of ringing the bells, and a corner requires a ringer to change the speed of their bell, which involves extra physical effort. Also related is the problem of decomposing a permutation into the minimum number of block transpositions – see [7].

Tangles appear naturally as a sub-problem in the context of graph-drawing, and this was our original motivation for the problems considered here. In order to simplify a visualization of a large graph, it is sometimes advantageous to “bundle” sets of nearby edges together [15, 16]. Since the edges may be required to appear in different orders at the two ends of a bundle, they must be permuted along its length, and it is desirable to do this in a helpful and visually appealing way. Paths with few moves (or few corners) tend to be easy to follow.

With practical applications in mind, it is worth noting that the tangles resulting from our constructions can often be improved slightly by local modifications. For example, in Figure 2(a), one may eliminate the two



(a) Bubble sort variant: use one R-move to route each path to its correct position, starting from the rightmost, $\pi(n)$. Path i may have $\Omega(i)$ L-moves.

(b) Odd-even sort: at alternate steps, apply swaps in all odd positions, or all even positions, wherever the two elements form an inversion.

FIGURE 3. Tangles constructed according to two natural greedy algorithms. Both require $\Omega(n^2)$ moves in the worst case as $n \rightarrow \infty$.

swaps where the tail and body of the “fish” meet, reducing the depth; in Figure 2(b), the isolated swap in the middle of the leftmost column may be moved upward to meet the swaps at the top, eliminating a move. Such modifications may be iterated, but will not improve the worst case asymptotic performance of the constructions.

Further notation and conventions. As mentioned above, it is convenient to consider a tangle in terms of its swaps, and we think of the swaps as located at elements of the integer lattice \mathbb{Z}^2 . If $\pi_t, \pi_{t+1} \in S_n$ are two consecutive permutations in a tangle, and they are related by non-overlapping

swaps thus: $\pi_t \cdot s(p_1) \cdots s(p_k) = \pi_{t+1}$, then we say that the tangle has swaps at **locations** $(p_1, t), \dots, (p_k, t)$. The first coordinate is sometimes called position, and increases from left (**West**) to right (**East**) (from 1 to $n - 1$); the second coordinate is called time, and increases from top (**North**) to bottom (**South**). If a tangle consists of permutations in S_n then we sometimes call n the **width** of the tangle.

We identify two tangles if they have the same set of swap locations; thus, we consider the tangle with permutations π_1, \dots, π_t to be the same as that with permutations $\gamma \cdot \pi_1, \dots, \gamma \cdot \pi_t$, for any permutation γ . In particular, a tangle that performs a permutation π may be equivalently be considered as starting at π^{-1} and ending at id , thus “sorting” π^{-1} . The latter convention was adopted in [4]. It will also be useful to allow times of swaps to take *any* value in \mathbb{Z} , and to identify two tangles if one is obtained from the other by adding a constant to all swap times (thus translating it vertically).

As mentioned earlier, we will construct tangles by combining smaller tangles (called gadgets), and for this it will be useful to translate horizontally as well as vertically. Thus, let $m < n$ and suppose that T is a tangle performing $\pi \in S_m$, with its swaps at locations $S \subset [1, m - 1] \times \mathbb{Z}$. Then for integers a, b , we may form a tangle T' of size n by placing swaps at the translated locations $S' := \{(i + a, t + b) : (i, t) \in S\}$; this performs the permutation $[1, \dots, a, \pi(a + 1), \dots, \pi(a + m), a + m + 1, \dots, n]$. Moreover, we may combine several tangles by taking the union of their sets of swap locations (perhaps after applying various translations).

A swap location (x, t) is called **even** or **odd** according to whether $x + t$ is even or odd. All the tangles we construct will have their swaps restricted to locations of one parity. As indicated above, a convenient way to highlight the structure of such a tangle is to draw a shaded 45° -rotated square centered at each swap, as in Figure 1(c). Recall that a move is a maximal non-vertical segment of a path. We call it an **L-move** if it runs in the North-East to South-West direction, and an **R-move** if it runs North-West to South-East.

Pattern-avoiding permutations will play a key role. (See e.g. [11] for background.) A **pattern** is a permutation $p \in S_m$. For $n \geq m$, we say that a permutation $\pi \in S_n$ (or, more generally, a sequence of n distinct real numbers π) **contains** the pattern p if there exist indices $1 \leq i_1 < \dots < i_m \leq n$ such that, for all $1 \leq j < k \leq m$, we have $\pi(i_j) < \pi(i_k)$ if and only if $p(j) < p(k)$. If π does not contain p then π is said to be **p -avoiding**.

The following concept will also be useful. For positive integers a_1, \dots, a_k with sum n , consider the partition of $[1, n]$ into the intervals $[1, a_1], [a_1 + 1, a_1 + a_2], \dots, [n - a_k + 1, n]$ with these lengths. We say that a permutation $\pi \in S_n$ is (a_1, \dots, a_k) -**split** if it maps each of these intervals to itself.

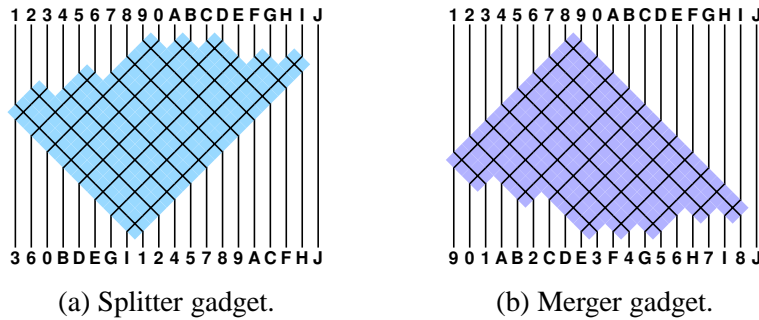


FIGURE 4

Organization of the paper. In Section 2 below we introduce the gadgets that will be used in our constructions, and prove their required properties. Proposition 4 and Theorems 1 and 2 are then proved in Sections 3–5 respectively. We prove the bound Theorem 3 in Section 6, via a combinatorial argument. In contrast, the lower bound in Proposition 5 and the fact that some permutations require 3 moves in some path (Proposition 17) are proved by explicitly exhibiting suitable permutations, in Sections 7 and 8 respectively. Although the permutations in question are very easy to describe, the proofs of both results are surprisingly delicate.

2. GADGETS

In this section we introduce the gadgets that will be used to prove Theorems 1 and 2. They come in three main categories, with several variants in each.

2.1. Splitter and Merger. Our first gadget comes in two variant forms, which are reflections of each other about a horizontal axis. A **splitter** gadget has swaps at locations

$$(i - j + a, -i - j),$$

for all $j \geq 0$ and $0 \leq i \leq b(j)$, where a is an even integer, and b is a non-decreasing integer-valued function of bounded support. Thus, a splitter consists of swaps at all even locations in a region bounded below by two line segments running South-East and North-East, and bounded above by an interface comprising any sequence of North-East and South-East segments. See Figure 4(a) for an example. The idea is that it separates the paths into two arbitrary sets, and places them on the left and right sides while maintaining the relative order within each set.

To formalize this: a permutation $\pi = [\pi(1), \dots, \pi(n)]$ is called **Grassmannian** if it has at most one descent, i.e. at most one index k such that $\pi(k) > \pi(k + 1)$.




FIGURE 5. A direct tangle, performing a 321-avoiding permutation.

Lemma 6. *A permutation can be performed by some splitter if and only if it is Grassmannian. Furthermore, the correspondence between splitters and Grassmannian permutations is bijective.*

For the purposes of the claimed bijectivity, recall that two tangles are identified if they have the same set of swap locations.

Proof of Lemma 6. The identity permutation is clearly performed by the trivial tangle containing no swaps. Any other Grassmannian permutation π has exactly one descent; say $\pi(k) > \pi(k+1)$. We take $a = k$, and $b(i) = \pi(k-i) - (k-i)$ for $0 \leq i \leq k-1$, and $b(i) = 0$ for $i \geq k$. The function b is easily seen to be non-decreasing. The lower boundary of the splitter consists of k steps South-East followed by $n-k$ steps North-East. The upper boundary also consists of k South-East steps and $n-k$ North-East steps, with the $\pi(i)$ th step being South-East if and only if $i \leq k$. For $i \leq k$, path $\pi(i)$ makes one L-move, starting at position i and ending at position $\pi(i)$. For $i > k$, path $\pi(i)$ similarly makes one R-move. The paths $\pi(i)$ for $i \leq k$ maintain their order relative to each other, as do those for $i > k$. See Figure 4(a). By similar reasoning, any splitter performs a Grassmannian permutation. Since different splitters perform different permutations, the correspondence is bijective. \square

A **merger** is obtained by reflecting a splitter about a horizontal axis. Thus it has swaps at all locations

$$(i - j + a, i + j),$$

for a and $b(\cdot)$ as above. See Figure 4(b). The corresponding permutation is the inverse of that performed by the splitter. A permutation π is the inverse of a Grassmannian permutation if and only if, for some k , the values $1, \dots, k$ appear in increasing order in the sequence $\pi = [\pi(1), \dots, \pi(n)]$, and so do $k+1, \dots, n$. The proof of the following lemma is immediate.

Lemma 7. *A permutation can be performed by some merger if and only if its inverse is Grassmannian. Furthermore, this correspondence is bijective.*

Both splitters and mergers are special cases of a more general class of tangles considered in [4], called *direct tangles*. A direct tangle is one in

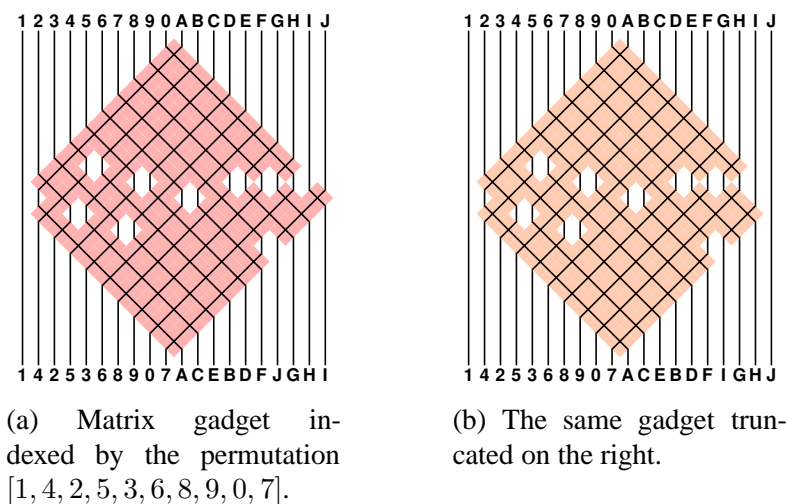


FIGURE 6

which each path has at most one move. Modulo a suitable convention regarding split permutations (in which the parts of the tangle corresponding to different splitting intervals may be translated vertically), such a tangle consists of swaps at all even locations within a region bounded above and below by interfaces comprising North-East and South-East segments. See Figure 5. It is shown in [4] that a permutation admits a direct tangle if and only if it is 321-avoiding. (Grassmannian permutations and their inverses are indeed 321-avoiding.) The correspondence is again bijective.

2.2. Matrix gadget. Let $n = 2m$ be even, and let $\alpha \in S_m$ be a permutation. The **matrix gadget** indexed by α consists of swaps at the locations

$$(i + j - 1, i - j)$$

for all pairs $i, j \in \{1, \dots, m\}$ except those with $\alpha(i) = j$. In other words, a square angled at 45° to the axes is filled with swaps at all odd locations, except for those locations corresponding to the support of the (rotated) permutation matrix of α . See Figure 6(a) for an example. The idea is that a matrix gadget performs any given permutation on one half; the following result says that the effect on the second half is the inverse permutation.

Lemma 8. *Let $n = 2m$ and let $\alpha \in S_m$. The matrix gadget indexed by α performs the permutation*

$$[\alpha(1), \alpha(2), \dots, \alpha(m), \alpha^{-1}(1)+m, \alpha^{-1}(2)+m, \dots, \alpha^{-1}(m)+m] \in S_n.$$

Proof. This is straightforward to check. Suppose $\alpha(i) = j$. Then path j makes an R-move until it encounters the “omitted swap” corresponding to




FIGURE 7. A pair of complementary paths in a matrix gadget. Horizontal positions are indicated along the top line.

the pair (i, j) , and then makes a vertical segment of length 1 followed by an L-move, finishing in position i . Similarly, path $i + m$ finishes in position $j + m$ after an L-move and an R-move. See Figure 7. \square

The matrix gadget is fundamentally more powerful than our other gadgets, in the sense that it can perform $(n/2)!$ different permutations in S_n , whereas each the others can only perform $O(c^n)$ permutations for some constants c . The matrix gadget is the source of the “holes” (or, more generally, clusters) mentioned in the introduction. Theorem 3 reflects the fact that some such construction is a requirement if we are to have only linearly many moves.

For some of our constructions, we will need the following variants of the matrix gadget for odd n . Let $n = 2m - 1$, and let $\alpha \in S_m$. The **truncated matrix gadget** indexed by α is simply the matrix gadget of the larger size $2m$ indexed by α , but with the rightmost swap (in location $(2m - 1, 0)$) omitted (whether or not it is present in the original matrix gadget). See Figure 6(b) for an example. This gadget performs a permutation of the form

$$[\alpha(1), \alpha(2), \dots, \alpha(m), \dots] \in S_{2m-1};$$

i.e. α on the m leftmost positions, and *some* permutation on the $m - 1$ rightmost positions. The precise nature of the permutation on the right will not matter for our applications. Similarly, we may truncate a matrix gadget on the left side to obtain any desired permutation on the rightmost m positions.

Finally, note the following subtle variation. If $n = 2m$ and the index permutation satisfies $\alpha(1) = 1$, then the standard matrix gadget *already* has no swap in the leftmost column. Figure 6(a) is an example. Therefore, it

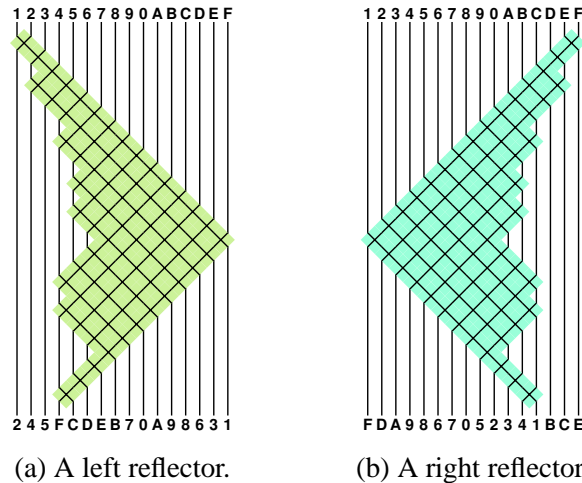


FIGURE 8

can also be regarded as a gadget involving only positions $2, \dots, 2m$, and performing any desired permutation on the positions $2, \dots, m$. (And it may then be translated one position leftward, for example).

2.3. Reflectors. Our final gadget also comes in two complementary forms, this time related by reflection in a vertical axis. A **right reflector** gadget consists of swaps at locations

$$(i + j + 1, j - i)$$

for all $j \geq 0$ and $0 \leq i \leq b(j)$, where b is a non-decreasing integer-valued function of bounded support. Thus, a right reflector consists of swaps at all odd locations in a region bounded on the left by two line segments running South-West and South-East, and bounded on the right by an interface comprising a sequence of South-West and South-East segments. See Figure 8(b). In this case, this rightmost bounding interface must stay to the left of the horizontal coordinate n . Therefore it corresponds to a Dyck path. The idea of the right reflector is that every path starts with an R-move, then has a vertical segment (now possibly of length greater than 1), and then is “reflected” back with an L-move.

Lemma 9. *A permutation can be performed by some right-reflector if and only if it is 132-avoiding. Furthermore, this correspondence between gadgets and permutations is bijective.*

Proof. We prove the “if” direction by induction on n . For $n = 1$, the claim is clear. For $n > 1$, suppose that π is 132-avoiding. Consider the location of element n in π , and write $\pi = [\alpha, n, \beta]$, where α, β are the sequences of




FIGURE 9. Inductive construction of a right reflector. The rectangle is chosen so as to route path n to its correct location, and the two remaining triangles are then filled with smaller right reflectors.

numbers to the left and right of n . Note that α and β are both 132-avoiding. Also, every element of α is greater than every element of β , otherwise we would have a 132 pattern including n .

We construct a right reflector as shown in Figure 9. There is a 45° rectangle filled with swaps, with one corner at $(1, 0)$ and an opposite corner at $(n-1, n-2\pi^{-1}(n))$; path n has a vertical segment until it hits this rectangle just above its rightmost corner, and then has an L-move. (A trivial case is when $\pi^{-1}(n) = n$, the rectangle is empty, and path n is vertical). Finally, we use the inductive hypothesis to insert two strictly smaller right reflectors, which perform the permutations corresponding to relative orders of α and β , in the triangular regions to the North-East and South-East of the rectangle.

We now turn to the “only if” direction. Suppose that a right reflector gadget T performs a permutation π . We first note that T is simple. Indeed, every path consists of an R-move, then a vertical segment, then an L-move (where it is possible that one or both of these moves is empty); since two paths can only cross during the R-move of one and the L-move of the other, they cannot cross more than once. Now suppose for a contradiction that π contains a 132 pattern. Thus, there exist $u < v < w$ with $\pi(u) < \pi(w) <$

$\pi(v)$. Consider the location x of the unique swap between paths $\pi(v)$ and $\pi(w)$. By the definition of the right reflector, every odd location in the 45° rectangle with corners $(1, 0)$ and x contains a swap. However, path $\pi(u)$ starts to the left of path $\pi(w)$, and traverses the entire rectangle during its R-move, and crosses path $\pi(v)$ at the South-East side of the rectangle. This contradicts simplicity.

To check bijectivity, since clearly every gadget performs only one permutation, it is enough to check that the two sets have equal cardinality. The number of 132-avoiding permutations in S_n is given by the Catalan number C_n . A right reflector gadget is encoded by a Dyck path describing its right boundary. Therefore the number of them is also C_n . See e.g. [17, Ex. 6.19]. \square

We remark that the standard Catalan recurrence $C_{n+1} = \sum_{i=0}^n C_i C_{n-i}$ is implicit in our inductive construction above. Arguments similar to ours appear in the context of *stack sorting* (see [19, p. 14] and [13]).

A **left reflector** gadget is simply the image of a right reflector under the reflection in the vertical line through the center of the permutation. Thus it has swaps at locations

$$(n - i - j, j - i)$$

for i, j and $b(\cdot)$ as before. See Figure 8(a). The next result follows immediately by symmetry.

Lemma 10. *A permutation can be performed by some left reflector if and only if it is 213-avoiding. This correspondence is bijective.*

In our applications, we will prove and use two properties of 132-avoiding (or 213-avoiding) permutations that are interesting in their own right: (i) any permutation can be decomposed into a cyclic permutation and a 132-avoiding permutation (Section 4); (ii) a 132-avoiding permutation can be found that maps any given subset of $\{1, \dots, n\}$ to any other subset of the same size (Section 5).

3. LOGARITHMIC MOVES PER PATH

Our simplest construction uses only splitters to obtain a simple tangle with logarithmically many moves per path.

Proof of Proposition 4. See Figure 10 for the construction and Figure 2(c) for an example. Let $\pi \in S_n$ be any permutation and let $m = \lfloor n/2 \rfloor$. Let $L = \{\pi(1), \dots, \pi(m)\}$ and $R = \{\pi(m+1), \dots, \pi(n)\}$. Consider the Grassmannian permutation ρ obtained by writing the elements of L in increasing order followed by the elements of R in increasing order. By Lemma 6 there is a splitter that performs ρ . We first apply this splitter. It




FIGURE 10. Construction for Proposition 4. After the initial splitter, the two rectangles signify smaller recursively defined versions of the same construction.

remains to perform $\rho^{-1} \cdot \pi$, which is an $(m, n - m)$ -split permutation. Thus, we can split into two subproblems. We then recursively apply the same procedure to each, and place the resulting tangles below the initial splitter, after appropriate translations.

Each path performs at most one move within each splitter that it encounters (perhaps fewer, since some may splitters involve no move for the path, and some pairs of splitters may be positioned to abut one another, so that two moves coalesce). A path encounters at most $\lceil \log_2 n \rceil$ splitters.

The tangle is simple, since if two paths cross in the first splitter, then they subsequently remain in the two distinct halves. \square

We remark that the above construction can be modified to obtain a tangle with only one cluster, and $O(\log n)$ moves per path, thus matching up to constants the extremal case $\theta \nearrow \frac{1}{2}$ of Theorem 3. After the first splitter, route path $\pi(m)$ alongside the South-West boundary of the splitter to its correct position m . This path then remains vertical for the rest of the tangle, keeping the two halves apart and preventing formation of holes. Iterate on the two intervals $[1, m - 1]$ and $[m + 1, n]$, and ensure that the subsequent splitters are translated upward until they touch some swap of a previous stage.

4. BOUNDED MOVES PER PATH

In this section we prove Theorem 1. The construction will make essential use of reflector gadgets. We use the following key property of 312-avoiding permutations, which we will then extend to other patterns of length 3. A permutation is called **cyclic** if it has only one cycle (or orbit).

Lemma 11. *For any permutation $\pi \in S_n$, there exists a 312-avoiding permutation σ such that $\sigma \cdot \pi$ is cyclic.*

Proof. Assume $n \geq 2$, otherwise the result is trivial. We use an iterative procedure to compute a suitable σ . We start with π , and pre-compose it by a sequence of suitably chosen disjoint cycles. The composition of these cycles will be 312-avoiding. Given the current permutation τ (which is initially equal to π), a **rainbow interval** is an interval $[a, b]$ such that all elements $i \in [a, b]$ belong to distinct cycles of τ . A **maximal** rainbow interval $[a, b]$ is one that is not a proper subset of another; thus, either we have $a = 1$, or $a - 1$ belongs to the same cycle as some element of $[a, b]$; a similar condition holds at the other end. If τ is not cyclic, then there exists some maximal rainbow interval $[a, b]$ of length at least 2. We now replace τ with the permutation $\tau' := \kappa \cdot \tau$, where

$$\kappa := \left[1, \dots, a-1, \underbrace{a+1, a+2, \dots, b, a}, b+1, \dots, n \right],$$

(i.e. a rotation of the interval $[a, b]$; note that κ is 312-avoiding). The effect of this change is to unite all the distinct cycles of the elements of $[a, b]$ into one cycle; all other cycles are unchanged. Consequently, if we iterate this operation, the rainbow intervals used at successive steps will be disjoint, and eventually τ will be cyclic. Moreover, the various cycles κ used at different steps commute with each other, and their composition σ is 312-avoiding. \square

Corollary 12. *For any permutation $\pi \in S_n$, and any pattern $p \in \{312, 231, 213, 132\}$, there exists a p -avoiding permutation σ such that $\sigma \cdot \pi$ is cyclic.*

Proof. Lemma 11 is the case $p = 312$. Let $\text{rev} := [n, n-1, \dots, 1] \in S_n$ be the reverse permutation. For the case $p = 231$, apply Lemma 11 to the conjugate permutation $\text{rev} \cdot \pi \cdot \text{rev}^{-1}$ to obtain a 312-avoiding σ with $\sigma \cdot \text{rev} \cdot \pi \cdot \text{rev}^{-1}$ cyclic. The conjugate of the last permutation by rev^{-1} is $\text{rev}^{-1} \cdot \sigma \cdot \text{rev} \cdot \pi \cdot \text{rev}^{-1} \cdot \text{rev} = (\text{rev}^{-1} \cdot \sigma \cdot \text{rev}) \cdot \pi$, which thus is cyclic also. The permutation $\text{rev}^{-1} \cdot \sigma \cdot \text{rev}$ is 231-avoiding, as required.

For $p = 213$, apply Lemma 11 to $\text{rev} \cdot \pi$, to obtain a 312-avoiding σ with $\sigma \cdot \text{rev} \cdot \pi$ cyclic. Then $\sigma \cdot \text{rev}$ is 213-avoiding. Finally, for $p = 132$, apply the conjugation trick to the $p = 213$ case. \square

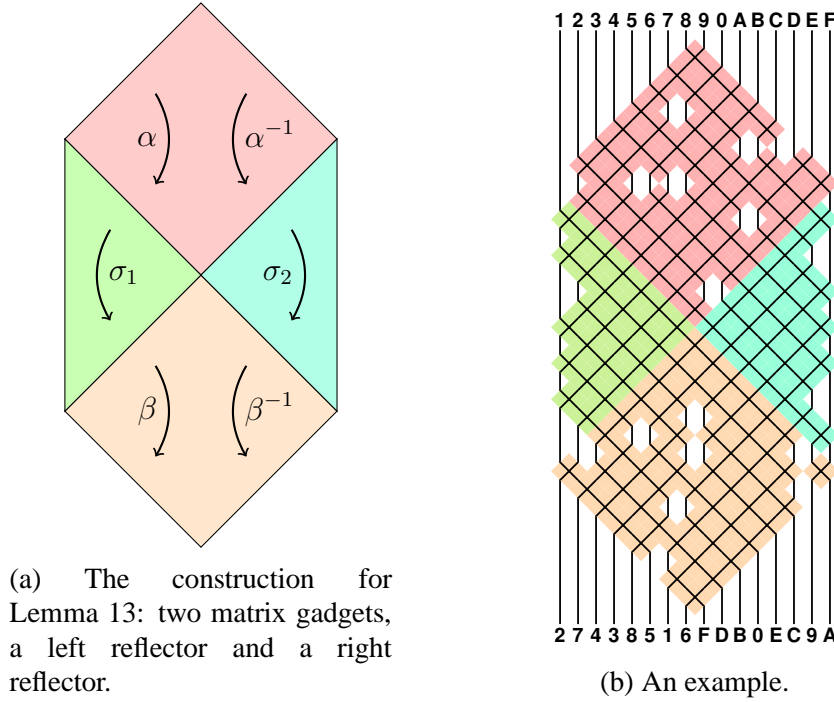


FIGURE 11

Here is the main step in the proof of Theorem 1.

Lemma 13. *For n even and any $(n/2, n/2)$ -split permutation $\pi \in S_n$, there is a tangle with at most 4 moves per path that performs π .*

Proof of Lemma 13. Let $n = 2m$. Since π is (m, m) -split, there exist $\pi_1, \pi_2 \in S_m$ such that $\pi_1 = (\pi(1), \dots, \pi(m))$ and $\pi_2 = (\pi(m+1) - m, \dots, \pi(2m) - m)$. We construct the required tangle using two matrix gadgets, one above the other, together with a left reflector and a right reflector (each of width m) in the two spaces between them, as in Figure 11.

Let $\alpha, \beta \in S_m$ be the permutations indexing the upper and lower matrix gadgets respectively (see the definition of a matrix gadget). Let $\rho_1, \rho_2 \in S_m$ be the permutations performed by the left reflector and the right reflector respectively. Clearly such a tangle performs an (m, m) -split permutation, for any choices of $\alpha, \beta, \rho_1, \rho_2$. Our task is to choose these permutations so as to perform the required π .

Recall from Lemma 8 that a matrix gadget performs its indexing permutation on the left and the inverse permutation on the right. Thus, our tangle performs π if and only if

$$(1) \quad \alpha \cdot \rho_1 \cdot \beta = \pi_1 \quad \text{and} \quad \alpha^{-1} \cdot \rho_2 \cdot \beta^{-1} = \pi_2.$$

The first equation gives $\pi_1^{-1} \cdot \alpha \cdot \rho_1 = \beta^{-1}$, and substituting into the second gives $\alpha^{-1} \cdot \rho_2 \cdot \pi_1^{-1} \cdot \alpha \cdot \rho_1 = \pi_2$. Rearranging,

$$(2) \quad \rho_2 \cdot \pi_1^{-1} = \alpha \cdot (\pi_2 \cdot \rho_1^{-1}) \cdot \alpha^{-1}.$$

There exists an α satisfying (2) if and only if the two permutations $\rho_2 \cdot \pi_1^{-1}$ and $\pi_2 \cdot \rho_1^{-1}$ are conjugate. By Corollary 12, for any π_1 , we can choose a 132-avoiding ρ_2 such that $\rho_2 \cdot \pi_1^{-1}$ is cyclic. Similarly, for any π_2 , we can choose a 213-avoiding ρ_1 such that $\rho_1 \cdot \pi_2^{-1}$ is cyclic, whence the inverse $\pi_2 \cdot \rho_1^{-1}$ is cyclic also. The permutations ρ_1, ρ_2 can be performed by the appropriate reflector gadgets by Lemmas 9 and 10. Thus, the two permutations mentioned above are both cyclic, and therefore conjugate, and so we can choose α satisfying (2). Finally, we can compute $\beta = \rho_1^{-1} \cdot \alpha^{-1} \cdot \pi_1$, and (1) will be satisfied.

The resulting tangle has at most 4 moves per path: a path has two moves in each matrix gadget, and these moves continue into the reflectors, since the gadgets abut each other. \square

Proof of Theorem 1. First consider even $n = 2m$. See Figure 2(a) for an example. Using Lemma 6, we first apply a splitter gadget that performs the permutation τ , where $\tau(1), \dots, \tau(m)$ are $\pi(1), \dots, \pi(m)$ in increasing order, and $\tau(m+1), \dots, \tau(2m)$ are $\pi(m+1), \dots, \pi(2m)$ in increasing order. We then use Lemma 13 to obtain a tangle that performs the (m, m) -split permutation $\tau^{-1} \cdot \pi$, and we place this tangle below the splitter. The splitter adds at most one move to each path.

For odd $n = 2m + 1$, we modify the construction as shown in Figure 12. The initial splitter separates the paths into sets of sizes m and $m + 1$, with path at the extreme right being $z = \max\{\pi(m+1), \dots, \pi(2m+1)\}$. We then proceed as before for the first $2m$ paths. Finally, we insert path z into its proper place in π by an L-move alongside the South-East side of the lower matrix gadget. Path z has only 2 moves, and every other path still has at most 5 moves. \square

We remark that the last trick for adding an additional path with only 2 moves could be iterated, to obtain an inductive construction of larger tangles. However, in general this would incur a quadratic number of moves in total, for similar reasons to the construction in Figure 3(a).

5. LINEAR TOTAL MOVES

We begin with another fact about 132-avoiding permutations. Write $[k] := \{1, \dots, k\}$.

Lemma 14. *If $A, B \subseteq [n]$ have equal cardinality then there exists a 132-avoiding $\pi \in S_n$ with $\pi(A) = B$.*




FIGURE 12. The construction for Theorem 1 for odd n : the largest element in the right half of the permutation is routed along the right side.

Proof. The proof is by induction on n . If $n = 1$ then the claim is obvious. Suppose that the theorem holds for all $n' < n$. We will deduce it for n . A pair (i, j) is called **conforming** if either $i \in A$ and $j \in B$, or $i \notin A$ and $j \notin B$. (In other words, if we are allowed to assign $\pi(i) = j$). We consider several cases.

Case 1. Pair $(n, 1)$ is conforming. Without loss of generality, suppose that $n \notin A$ and $1 \notin B$; otherwise take complements of A and B . Consider the set $B - 1 := \{i - 1 : i \in B\}$. By the induction hypothesis, there exists a 132-avoiding $\sigma \in S_{n-1}$ with $\sigma(A) = B - 1$. Define $\pi \in S_n$ by setting $\pi(n) = 1$, and $\pi(i) = \sigma(i) + 1$ for $i < n$. Then π is 132-avoiding, and maps A to B , as required.

Case 2. Pair $(1, n)$ is conforming. Without loss of generality, $1 \notin A$ and $n \notin B$. Consider $A - 1 := \{i - 1 : i \in A\}$. By the induction hypothesis there exists a 132-avoiding $\sigma \in S_{n-1}$ with $\sigma(A - 1) = B$. Define $\pi \in S_n$ by $\pi(1) = n$ and $\pi(i) = \sigma(i - 1) + 1$ for $i > 1$.

Case 3. Pair (n, n) is conforming. Apply the inductive hypothesis to $1, \dots, n - 1$ and set $\pi(n) = n$.

Case 4. None of the pairs $(n, 1)$, $(1, n)$, (n, n) is conforming. Without loss of generality, $1 \in A$. Then $n \notin B$ because $(1, n)$ is not conforming. Then $n \in A$ because (n, n) is not conforming. Then $1 \notin B$ because $(n, 1)$ is not conforming. In summary, we have $1, n \in A$ but $1, n \notin B$.

We claim that there exists an integer k with $2 \leq k \leq n - 2$ such that $|A \cap [k]| = |B \cap ([n] \setminus [n - k])|$. Indeed, we have $|A \cap [1]| = 1 > 0 = |B \cap ([n] \setminus [n - 1])|$, whereas $|A \cap [n - 1]| = |A| - 1 < |B| = |B \cap ([n] \setminus [1])|$; but the difference $|A \cap [j]| - |B \cap ([n] \setminus [n - j])|$ decreases by at most 1 as j is increased by 1; thus it must be 0 for some j .

Let $A' = A \cap [k]$ and $B' = (B \cap ([n] \setminus [n - k])) - (n - k)$. By the induction hypothesis, (since $k < n$) there exists a 132-avoiding $\pi_1 \in S_k$ with $\pi_1(A') = B'$. Let $A'' = (A \cap ([n] \setminus [k])) - i$ and $B'' = B \cap [n - k]$. By the induction hypothesis, (since $n - k < n$) there exists a 132-avoiding $\pi_2 \in S_{n-k}$ with $\pi_2(A'') = B''$. We define π by setting $\pi(j) = \pi_1(j) + n - k$ for $j \leq k$, and $\pi(j) = \pi_2(j - k)$ for $j > k$. This π is 132-avoiding: if $u < v < w$ form a 132 pattern, then we cannot have all three in $[k]$ or all three in $[n] \setminus [k]$. On the other hand, we cannot have $u \leq k < w$: indeed, for all $i \leq k < j$ we have $\pi(i) > n - k \geq \pi(j)$. \square

The following is a major ingredient of the proof of Theorem 2.

Proposition 15. *Let $\pi \in S_n$ be a $(\lceil n/2 \rceil, \lfloor n/2 \rfloor)$ -split permutation. The permutation π can be performed by a tangle all of whose swaps are within the triangular region $\{(x, t) : -x < t < x\}$. The tangle accepts n paths running in the South-East direction on its North-West edge, and outputs them running in the South-East direction on its South-East edge, and has at most $4n$ moves including these input and output segments.*

Proof. We first assume that $n = 2m$ is even, so π is (m, m) -split. The construction of the required tangle C is recursive: it consists of a matrix gadget M , together with a right reflector R of width m placed to the North-East of the matrix, and a smaller, recursively-constructed version C' of itself (performing a suitable permutation of size m) placed to the South-East of the matrix. See Figure 13(a).

We now explain how to choose the gadgets. Let $\rho, \mu \in S_n$ be the permutations performed by the right reflector R (when translated to the right half $[m + 1, n]$) and the matrix gadget M , respectively. Since the right-reflector does not affect positions in the left half $[1, m]$, we require that $[\mu(1), \dots, \mu(m)] = [\pi(1), \dots, \pi(m)] (\in S_m)$. Therefore we choose the matrix gadget to be indexed by this last permutation. Now consider the right half. The tangle C' can perform any desired $(\lceil m/2 \rceil, \lfloor m/2 \rfloor)$ -split permutation on positions $m + 1, \dots, 2m$. Therefore, letting $Q = [m + \lceil m/2 \rceil + 1, n]$ be the set of positions in the last quarter of $[1, n]$, we need to choose ρ so that $\rho \cdot \mu(Q) = \pi(Q)$. Since $|\mu(Q)| = |\pi(Q)|$, by Lemmas 9 and 14, there is a right reflector that achieves this.

In the case when $n = 2m + 1$ is odd, the construction is modified as follows. The matrix gadget is replaced with a truncated version (with the

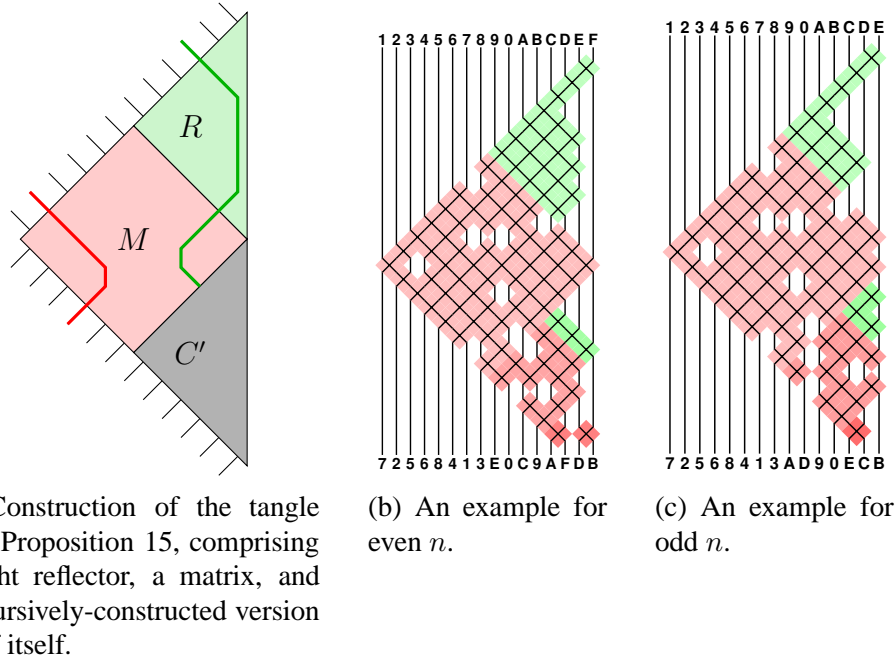


FIGURE 13

rightmost swap deleted), so that we may choose it to perform the required permutation on positions $1, \dots, m + 1$.

Finally, we count moves. Suppose that all paths start running in the South-East direction. Then each path makes at most 2 moves in the reflector together with the matrix, including the input path, but not including the final R-move in the case of the paths in the right half. Since these moves continue into C' , writing $A(n)$ for the maximum number of moves required by our construction for a permutation of size n , we have

$$A(n) \leq 2n + A(\lfloor n/2 \rfloor).$$

By induction, $A(n) \leq 4n$. \square

Proof of Theorem 2. The construction is illustrated in Figure 14, and Figure 2(b) is an example. We first assume that n is a multiple of 4, and write $n = 4m$. As shown in Figure 14, the tangle finishes with a merger G that (by Lemma 7) intersperses the paths in locations $1, \dots, 2m$ with those in $2m + 1, \dots, 4m$ in an arbitrary way while maintaining the relative order of each. Therefore, the remainder of the tangle (above the merger) needs to perform an arbitrary $(2m, 2m)$ -split permutation. On the other hand, the tangle starts with two splitters S_1 and S_2 , placed in the first and second halves. By Lemma 6 each of these splitters can map any desired set of

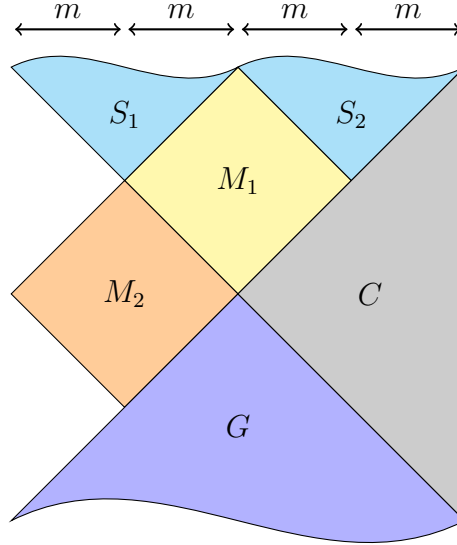


FIGURE 14. Construction for Theorem 2: splitters S_1, S_2 , matrix gadgets M_1, M_2 , merger G , and a tangle C from Proposition 15.

paths into its own first half (of width m). Therefore, the task for the remaining portion of the tangle (i.e. everything apart from the merger and the two splitters) is to perform an arbitrary (m, m, m, m) -split permutation.

The remainder of the tangle is composed of two matrix gadgets, together with a tangle constructed via Proposition 15. Both matrix gadgets have width $2m$. The upper matrix gadget M_1 occupies the middle half $[m + 1, 3m]$ of $[1, n]$. The other matrix gadget, M_2 , abuts M_1 to the South-West and occupies the first half $[1, 2m]$. The tangle C from Proposition 15 also has width $2m$, and is located on the right, partially abutting M_1 .

We now explain how to choose these gadgets. The matrix gadget M_2 is chosen so as to perform the required permutation in the first quarter $[1, m]$. Then M_1 chosen so that the required permutation in the second quarter $[m + 1, 2m]$ is performed by the left half of M_1 composed with the right half of M_2 . Finally, C needs to perform an arbitrary (m, m) -split permutation (on positions $[2m + 1, 4m]$). This can be achieved, by Proposition 15.

We now count moves. We first total the moves within each component. When two components abut each along a common boundary, the moves crossing this boundary will be double-counted. Therefore we then subtract a term corresponding to the total length of the common boundaries. The

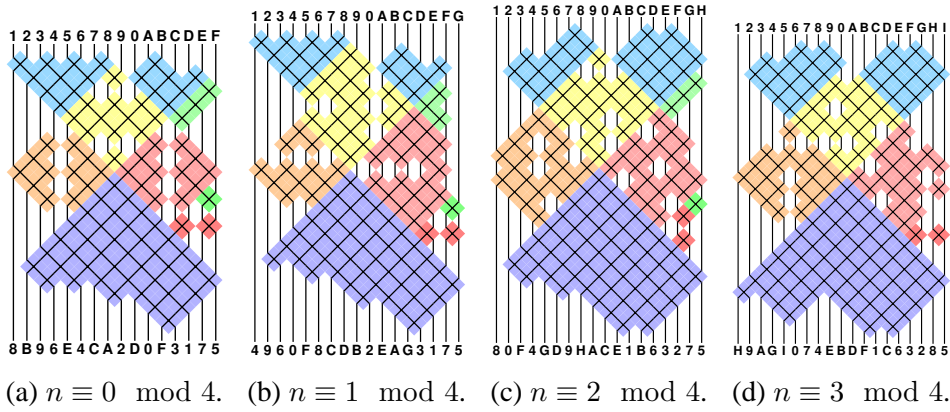


FIGURE 15. Variations of the construction for Theorem 2, according to the congruence class of n .

upper splitters each contribute $2m$ moves; the two matrix gadgets each contribute $4m$ moves; the final merger contributes $4m$ moves; and the tangle C contributes $4(n/2) = 2n$ moves, by Proposition 15. The total over-counting from common boundaries is $m + m + 3m + 3m$. Therefore, there are at most $24m - 8m = 16m = 4n$ moves.

Finally, we describe how the construction is adjusted when n is not a multiple of 4. Let $n = 4m + r$ where m is an integer and $r \in \{0, 1, 2, 3\}$. Depending on the value of r , we choose a suitable splitting into quarters, and use carefully chosen truncated matrix gadgets. The splitters and merger are adjusted to that the remaining central section of the tangle must perform a permutation that is split as follows:

$$\begin{aligned}
 r = 0 : & \quad (m, \quad m, \quad m, \quad m) \\
 r = 1 : & \quad (m, \quad m, \quad m + 1, \quad m) \\
 r = 2 : & \quad (m, \quad m + 1, \quad m + 1, \quad m) \\
 r = 3 : & \quad (m + 1, \quad m + 1, \quad m + 1, \quad m).
 \end{aligned}$$

The case $r = 0$ was described above. In the case $r = 1$, the matrix gadget M_1 is not truncated, but has width $2(m + 1)$, and is chosen to have no swap in its leftmost column. In the case $r = 2$, the matrix M_2 is truncated on its left side. In the case $r = 3$, both matrices have width $2(m + 1)$, and neither is truncated. For each of $r = 1, 2, 3$, the tangle C has odd width, and performs a $(m + 1, m)$ -split permutation, as stated in Proposition 15. These choices ensure that the various components can still abut each other without introducing extra moves at the boundaries. See Figure 15 for examples. \square

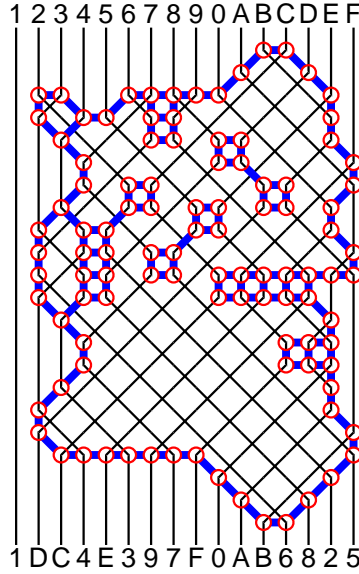


FIGURE 16. Corners and clusters (for a tangle constructed according to the proof of Theorem 2). Corners are circled, and corners connected by thick lines belong to the same cluster. There are three clusters.

We remark that, in the above construction, while the average number of moves per path is only 4, some paths may have as many as $\Theta(\log n)$ moves – this is a consequence of the recursive construction in Proposition 15.

6. CLUSTER BOUND

In this section we prove Theorem 3. Recall that swaps are located at elements of the integer lattice \mathbb{Z}^2 , and thus corners are located at elements of $(\mathbb{Z} + \frac{1}{2})^2$. Recall that a cluster is a connected component of the graph whose vertices are corners, and with an edge between two corners if their locations are within ℓ^∞ -distance 1. See Figure 16 for an example.

We start with a standard estimate for counting clusters. Let \mathbb{Z}_*^2 be the graph with vertex set \mathbb{Z}^2 and an edge between any two elements that are at ℓ^∞ -distance 1 from each other. By a ***-animal** we mean a finite subset of \mathbb{Z}^2 that induces a connected subgraph of \mathbb{Z}_*^2 . The **size** of a *-animal is the number of its vertices. Two *-animals are said to be **equivalent** if one can be obtained from the other by a translation of \mathbb{Z}^2 .

Lemma 16. *The number of equivalence classes of *-animals of size m is at most A^m , where $A = 7^7/6^6$.*

Proof. Apply the argument of Eden [6], adapted to the $*$ lattice. See also e.g. [12]. \square

Proof of Theorem 3. Fix $\theta \in (0, \frac{1}{2})$ and $n > \theta^{-8/\theta}$. Suppose for a contradiction that there are at least $e^{-n}n!$ distinct permutations $\pi \in S_n$ each of which has a tangle T_π with fewer than $K := (\frac{1}{2} - \theta)n$ clusters and fewer than $C := \frac{1}{6}\theta n \log n$ corners.

For any tangle T , suppose that there are no corners at the time $t + \frac{1}{2}$. It is easy to check that all segments must be vertical at the point, so the three permutations corresponding to times $t - \frac{1}{2}, t + \frac{1}{2}, t + \frac{3}{2}$ are all equal. Therefore we can remove one of these permutations from the sequence to obtain a new tangle. This operation preserves the number of corners, and does not increase the number of clusters. We can therefore assume that each of the tangles T_π defined above has depth at most equal to its number of corners. We may further assume that the time of the first corner is $\frac{1}{2}$. Therefore all corners are within a fixed rectangle R of area Cn . (Recall that there are at most C corners).

If we are given the set of locations of corners of a tangle, together with the directions of the two incident path segments at each corner, then we can recover the tangle. At any given corner there are at most $3^2 - 3 = 6$ possible choices for this pair of directions.

We now bound from above the number of possible tangles T_π . A cluster of size m corresponds to a $*$ -animal together with a location in the rectangle R . Therefore the number of possible tangles is at most

$$\sum_{m_1, \dots, m_k} \prod_{i=1}^k (A^{m_i} 6^{m_i} Cn),$$

where the sum is over all sequences $(m_i)_{i=1, \dots, k}$ with $k \leq K$, and $m_i \geq 1$ and $\sum_i m_i \leq C$, and where A is the constant from Lemma 16. The number of choices of such $(m_i)_{i=1, \dots, k}$ is at most 2^C , so the above expression is at most $(Cn)^K (12A)^C$.

Since each T_π corresponds to a different permutation π , we have

$$e^{-n}n! \leq (Cn)^K (12A)^C.$$

Taking logarithms, substituting for C and K , and using $\log(n!) > n \log n - n$, we obtain

$$n \log n - 2n \leq (\frac{1}{2} - \theta)n \log(Cn) + \frac{1}{6}\theta n \log n \log(12A).$$

Using $\log(12A) < 6$ and simplifying gives

$$\frac{1}{2} \log n - 2 \leq (\frac{1}{2} - \theta) \log C.$$

Since $0 < \frac{1}{2} - \theta < \frac{1}{2}$ and $\log C \leq \log n + \log \log n$, this implies

$$\theta \log n \leq 2 + \frac{1}{2} \log \log n.$$

It is straightforward to check that this gives a contradiction if $n > \theta^{-8/\theta}$. \square

We remark that there is nothing special about the choice of ℓ^∞ -distance 1 in the definition of clusters, except that it is fairly natural in the context of the tangles that we have constructed. The above argument goes through (with different constants) for other choices of norm and threshold distance.

7. LOWER BOUND FOR SIMPLE TANGLES

Proof of Proposition 5. First assume that $n = r^2 + 2$. Consider the permutation

$$(3) \quad \pi = \left[n, \underbrace{r+1, r, \dots, 2}, \underbrace{2r+1, 2r, \dots, r+2}, \dots, \underbrace{n-1, \dots, n-r}, 1 \right].$$

Thus, π consists of r blocks of length r , with each block having its elements in reverse order, and with 1 and n in reverse order at the two ends. For example, for $r = 3$ the permutation is $\pi = [11, 4, 3, 2, 7, 6, 5, 10, 9, 8, 1]$.

Define **block** i to be the set $B(i) = \{ir + 2, \dots, ir + r + 1\}$ for $i = 0, 1, \dots, r - 1$. Let T be a simple tangle that performs π . Every path other than 1 and n has at least two moves, since it crosses paths n and 1 in different directions. Observe that paths of $B(i)$ and $B(j)$ do not cross each other for $i \neq j$. Call a path **bad** if it has at least 3 moves, and call a block **terrible** if it contains at most one non-bad path. Next, we show that there are at most 3 non-terrible blocks, from which the result will follow easily.

Since paths from different blocks cannot cross each other, for any $i < j$, all elements of $B(i)$ precede all elements of $B(j)$ in any permutation of the tangle. Now consider the location (x, t) of the unique swap between paths 1 and n . Recall that (x, t) occurs between permutations π_t and π_{t+1} , and swaps the elements in locations x and $x + 1$. Let

$$H := \{\pi_t(x - r - 1), \dots, \pi_t(x + r)\}$$

be the set of elements that are within distance r on the left and right just before this swap. The set H has exactly $2r$ elements including 1 and n . Thus, by the previous observation, H contains elements from at most 3 blocks. We will show that any block having no elements in H is terrible.

Suppose that $B(i) \cap H = \emptyset$. By the argument of the previous paragraph, either all elements of $B(i)$ are before all elements of H in the permutation π_t , or they are after. Without loss of generality, assume the former. This




FIGURE 17. The key step in the proof of Proposition 5. The paths of H' must all cross path n before time s , therefore at least as many paths must cross path p during the same time interval.

implies that each path of $B(i)$ crosses 1 before it crosses n . Let $p < q$ be any elements of $B(i)$. We will show that at least one of paths p, q is bad. Let (y, s) be the location of the swap of p and q , and consider the permutation π_s . We consider six cases.

Suppose first that $\pi_s = [\dots, 1, \dots, p, q, \dots, n, \dots]$ (which is to say that p and q swap in the region above paths 1 and n). Path p has an R-move (to swap with q), then an L-move (to swap with 1), then an R-move (to swap with n). Therefore p is bad. The case $\pi_s = [\dots, n, \dots, p, q, \dots, 1, \dots]$ (where p and q swap below paths 1 and n) can be treated symmetrically.

Suppose now that $\pi_s = [\dots, p, q, \dots, 1, \dots, n, \dots]$ (which is to say that p and q swap in the region left of paths 1 and n , and at or before time t , so $s \leq t$). The argument for this case is illustrated in Figure 17. If p is not bad, then path p has an L-move (to swap with 1), followed by an R-move during which it swaps with both q and n . Let $H' := \{\pi_t(x - r - 1), \dots, \pi_t(x)\}$. All elements of H' are between p and n in π_t . These elements do not swap with p after time t , because $\pi_t(x) = 1$ has already swapped with p , while the others belong to different blocks and so never swap with p . Let u be the time of the swap of p and n . Since $u \geq t$, all elements of H' must swap with n strictly before time u . Therefore $u - t \geq r$. Therefore, path p has at least r swaps at times in the interval $[t, u)$ (since $s \leq t$, so its unique R-move is in progress throughout this interval). Since path p also swaps with 1 and n , it has at least $r + 2$ swaps in total, which contradicts simplicity, since p is involved in only $r + 1$ inversions. Thus, p is bad. The case $\pi_s = [\dots, p, q, \dots, n, \dots, 1, \dots]$ can be treated symmetrically.

Finally, the cases $\pi_s = [\dots, 1, \dots, n, \dots, p, q, \dots]$ and $\pi_s = [\dots, n, \dots, 1, \dots, p, q, \dots]$ are impossible, since together with our assumption about π_t they imply contradictions to simplicity.

Now we count moves. There at least $r - 3$ terrible blocks, each of which has at least $r - 1$ bad paths, which have at least 3 moves, so the total number of moves is at least

$$3(r - 3)(r - 1) \geq 3(r^2 - 4r) \geq 3n - c\sqrt{n}$$

for some $c > 0$.

For general n , we use the same construction with $r = \lfloor \sqrt{n - 2} \rfloor$, add an extra $n - r^2 - 2 < 2\sqrt{n} + 1$ elements at the end of the permutation, and adjust the constant. \square

It is tempting to try to extend the ideas of the above proof to show that there are permutations for which any simple tangle has $\gg n$ moves as $n \rightarrow \infty$ (perhaps even $\Omega(n \log n)$). A candidate permutation might be constructed recursively: a “level- k permutation” would have the same structure as π above, except with each block replaced with a smaller level- $(k - 1)$ permutation; the number of levels might be chosen to be of order $\log n$ (or at least something $\gg 1$). We have not succeeded in completing such an argument. Indeed, we do not know whether in fact $O(n)$ moves (or even $O(1)$ moves per path) suffice for a simple tangle.

8. PER PATH LOWER BOUND

Finally, we prove a lower bound on moves per path that applies even for non-simple tangles, as mentioned in the introduction.

Proposition 17. *For any $n > 8$ there exists a permutation $\pi \in S_n$ such that any tangle performing π has a path with at least 3 moves.*

Our proof of this seemingly simple statement is surprisingly intricate, and involves the two lemmas below. The permutation will be

$$\pi := [n, 3, 2, n - 3, n - 4, \dots, 5, 4, n - 1, n - 2, 1].$$

Recall that a pair of elements i, j is said to be an inversion of a permutation π if $i < j$ but $\pi^{-1}(i) > \pi^{-1}(j)$.

Lemma 18. *Let T be a tangle performing any permutation of the form $\pi = [n, \dots, 1]$ with each path making at most 2 moves. Let $1 < i < j < n$. If i, j is an inversion then paths i and j cross each other exactly once. If i, j is not an inversion then paths i and j either do not cross or cross exactly twice. In the latter case, the permutation at the time t just before paths 1 and n cross is of the form $\pi_t = [\dots, j, \dots, 1, n, \dots, i, \dots]$.*




FIGURE 18. Illustration of the proof of Lemma 19: the region formed by paths $a, b, 1, n$, and the four types of path that may intersect it. An RL path necessitates an LR path, but an LR path requires two further moves in order to cross paths n and 1 .

Proof. Paths i and j must cross an odd number of times if i, j is an inversion, and an even number of times if not. Since each path has at most 2 moves, they cannot intersect more than twice.

Suppose that paths i and j cross twice. Then i must have an R-move followed by an L-move, and vice-versa for j . Since any path other than 1 and n must cross path 1 during an L-move and cross path n during an R-move, the claimed form of π_t follows. \square

Lemma 19. *Let T be a tangle performing a permutation of the form $\pi = [n, \dots, 1]$ with each path making at most 2 moves. Let $z < a < b$ be some paths of T that first cross n and then cross 1 . If path z crosses neither a nor b , then a and b do not cross each other.*

Proof. Suppose on the contrary that paths a and b cross. By Lemma 18, they cross only once. Path b cannot cross a before n , since then b would have more than 3 moves. Similarly, path b cannot cross a after 1 , since a would have 3 moves.

Therefore, path b crosses n , then a , then 1 . Let N, E, S, W be the intersection points of the pairs of paths $(n, a), (a, b), (1, b), (1, n)$ respectively, all of which are unique by Lemma 18. These points are connected in clockwise order by four portions of the paths $a, b, 1, n$, which bound a region $NESW$. See Figure 18. Note that any path other than 1 or n has exactly one L-move and one R-move. Therefore, the sides NE and ES (which

form part of paths a and b) are straight line segments. On the other hand, the sides SW and WN may each contain at most one vertical segment, since paths 1 and n may have two moves in the same direction separated by a vertical segment.

Let $\ell(SW)$ denote the number of intersections of the side SW with paths other than 1, n , a , b (which corresponds to the length of its non-vertical portions), and similarly for each of the other three sides. By the above observations,

$$\ell(WN) \leq \ell(ES); \quad \ell(SW) \leq \ell(NE).$$

Every path other than 1, n , a , b than intersects NE must do so either in a single L-move or R-move, or with an L-move followed by an R-move, or vice-versa. Let $p(L), p(R), p(LR), p(RL)$ denote the numbers of paths in each category. We have

$$\begin{aligned} \ell(WN) &= p(RL) + p(R); & \ell(NE) &= p(LR) + p(L); \\ \ell(ES) &= p(LR) + p(R); & \ell(SW) &= p(RL) + p(L). \end{aligned}$$

Combining these equations with either of the above inequalities gives

$$p(RL) \leq p(LR).$$

We have $p(RL) \geq 1$, because of path z . However, $p(LR) \geq 1$ gives a contradiction, because such a path has at least 4 moves, in order to cross n , a , b and 1. \square

Proof of Proposition 17. Consider

$$\pi := \left[n, \underbrace{3, 2}, \underbrace{n-3, n-4, \dots, 5, 4}, \underbrace{n-1, n-2}, 1 \right].$$

We denote $A = \{2, 3\}$, $B = \{4, \dots, n-3\}$, and $C = \{n-2, n-1\}$.

Suppose for a contradiction that there exists a tangle T performing π in which each path has at most 2 moves. First suppose that T is simple. In each permutation of the tangle, the elements of A precede the elements of B , which precede the elements of C . Let t be the time of the swap 1, n , and suppose that some element x appears to the right of this swap, i.e. $\pi_t = [\dots, 1, n, \dots, x, \dots]$. If $x \in A \cup B$ then paths $x < n-1 < n-2$ contradict Lemma 19. Thus $x \in C$. Similarly, if $\pi_t = [\dots, y, \dots, 1, n, \dots]$ then $y \in A$. Thus there is no possible location for the elements of B in π_t , a contradiction.

Suppose on the other hand that T is not simple. Thus there exist paths i, j that **double-cross** (i.e. have two crossings). By Lemma 18, the pair i, j is not an inversion, therefore i, j are from two different sets among A, B, C . We claim that there exist $i' \in A$ and $j' \in C$ whose paths double-cross. Suppose not. Without loss of generality, assume that $i \in A$ and $j \in B$ double-cross. Since path i and any path of C do not double-cross, they do

not cross at all by Lemma 18. Since path i has an R-move then an L-move, we have $\pi_t = [\dots, 1, n, \dots, i, \dots]$. So paths $i < n - 1 < n - 2$ contradict Lemma 19. Thus i', j' exist as claimed.


Since $n > 8$ and $|B| > 2$, there are at least two elements u, v of B that either both cross 1 before n , or both cross n before 1. Without loss of generality, assume the latter. Since paths i' and u both move right then left, they cannot double-cross, and therefore by Lemma 18, they do not cross. By the same reasoning, i' and v do not cross. But now the paths $i' < u < v$ give a contradiction to Lemma 19. \square

ACKNOWLEDGEMENTS


We thank Omer Angel, Franz Brandenburg, David Eppstein, Martin Fink, Michael Kaufmann, Peter Winkler and Alexander Wolff for valuable conversations.

OPEN PROBLEMS

1. What is the asymptotic behavior as $n \rightarrow \infty$ of the maximum over permutations $\pi \in S_n$ of the minimum number of moves among *simple* tangles that perform π ? In particular, is it $O(n)$? Our results show only that it is between $3n - o(n)$ and $O(n \log n)$.
2. Similarly, what is the asymptotic behavior of the number of moves in the worst path (again, for the best simple tangle performing the worst permutation)? Our bounds are 3 and $O(\log n)$.
3. For general (not necessarily simple) tangles, what is the smallest constant a for which there exists a tangle with at most an moves for every permutation in S_n , for every n ? And what is the smallest b for which we can achieve at most b moves per path? We know that $2 \leq a \leq 4$ and $3 \leq b \leq 5$.
4. Many natural questions arise concerning permutations that can be performed by tangles of various restricted types. For example, suppose that the swaps of a tangle occupy all even locations in a simply connected region bounded above and below by interfaces consisting of North-East and South-East steps, and on the left and right by interfaces of South-West and South-East steps, as in Figure 19(a). Note in particular that there is one cluster, and no “holes”. It is not difficult to show that any permutation can be performed by such a tangle of depth at most $O(n^2)$ (see Figure 19(b) for the idea), but this seems far too large. What is the minimum depth needed? Is there a simple characterization of the set of permutations that can be performed if the depth is restricted to be at most n (say)?



(a) A tangle occupying a simply connected region bounded by monotone interfaces, as discussed in open problem 4.



(b) A greedy construction of such a tangle: we apply alternate rows of swaps in odd and even positions until path $\pi(n)$ is in the rightmost position, then continue in the same way with locations $1, \dots, n - 1$.

FIGURE 19

REFERENCES

- [1] M. Ajtai, J. Komlós, and E. Szemerédi. Sorting in $c \log n$ parallel steps. *Combinatorica*, 3(1):1–19, 1983.
- [2] M. H. Albert, R. E. L. Aldred, M. Atkinson, H. P. van Ditmarsch, C. C. Handley, D. A. Holton, and D. J. McCaughan. Compositions of pattern restricted sets of permutations. *Australian J. Combinatorics*, 37:43–56, 2007.
- [3] O. Angel, A. E. Holroyd, D. Romik, and B. Virag. Random sorting networks. *Advances in Mathematics*, 215(2):839–868, 2007.
- [4] S. Bereg, A. E. Holroyd, L. Nachmanson, and S. Pupyrev. Drawing permutations with few corners. In *Graph drawing*, volume 8242 of *Lecture Notes in Comput. Sci.*, pages 484–495. Springer, Cham, 2013.
- [5] N. Bergeron and S. Billey. RC-graphs and Schubert polynomials. *Experiment. Math.*, 2(4):257–269, 1993.
- [6] M. Eden. A two-dimensional growth process. In *Proc. 4th Berkeley Sympos. Math. Statist. and Prob., Vol. IV*, pages 223–239. Univ. California Press, Berkeley, Calif., 1961.

- [7] H. Eriksson, K. Eriksson, J. Karlander, L. Svensson, and J. Wästlund. Sorting a bridge hand. *Discrete Math.*, 241(1-3):289–300, 2001. Selected papers in honor of Helge Tverberg.
- [8] S. Felsner. On the number of arrangements of pseudolines. In *Proceedings of the 12th Annual Symposium on Computational Geometry*, pages 30–37, 1996.
- [9] S. Fomin and A. N. Kirillov. The Yang-Baxter equation, symmetric functions, and Schubert polynomials. *Discrete Mathematics*, 153(13):123 – 143, 1996. Proceedings of the 5th Conference on Formal Power Series and Algebraic Combinatorics.
- [10] A. M. Garsia. *The saga of reduced factorizations of elements of the symmetric group*. Université du Québec [Laboratoire de combinatoire et d’informatique mathématique (LACIM)], 2002.
- [11] S. Kitaev. *Patterns in permutations and words*. Monographs in Theoretical Computer Science. An EATCS Series. Springer, Heidelberg, 2011. With a foreword by Jeffrey B. Remmel.
- [12] D. A. Klarner. Cell growth problems. *Canad. J. Math.*, 19:851–863, 1967.
- [13] D. E. Knuth. *The art of computer programming. 1, Fundamental algorithms*. Addison-Wesley, 1973.
- [14] D. E. Knuth. *The art of computer programming. 3, Sorting and Searching*. Addison-Wesley, 1997.
- [15] S. Pupyrev, L. Nachmanson, S. Bereg, and A. E. Holroyd. Edge routing with ordered bundles. In *Graph Drawing*, pages 136–147, 2011.
- [16] S. Pupyrev, L. Nachmanson, and M. Kaufmann. Improving layered graph layouts with edge bundling. In *Proc. 18th Int. Symp. on Graph Drawing*, pages 465–479, 2010.
- [17] R. P. Stanley. *Enumerative Combinatorics*, volume 2. Cambridge Studies in Advanced Mathematics 62, Cambridge University Press, 1999.
- [18] D. C. Wang. Novel routing schemes for IC layout, part I: Two-layer channel routing. In *Proc. 28th ACM/IEEE Design Automation Conference*, pages 49–53, 1991.
- [19] J. West. *Permutations with forbidden subsequences, and, stack-sortable permutations*. PhD thesis, Dept. of Mathematics, Massachusetts Institute of Technology, 1990.
- [20] A. T. White. Ringing the cosets. *Amer. Math. Monthly*, 94(8):721–746, 1987.
- [21] B. Young. A Markov growth process for Macdonald’s distribution on reduced words, 2014. arXiv:1409.7714.

(Sergey Bereg) UNIVERSITY OF TEXAS AT DALLAS, USA

E-mail address: `besp@utdallas.edu`

(Alexander E. Holroyd) MICROSOFT RESEARCH, USA

E-mail address: `holroyd@microsoft.com`

(Lev Nachmanson) MICROSOFT RESEARCH, USA

E-mail address: `levnach@microsoft.com`

(Sergey Pupyrev) UNIVERSITY OF ARIZONA, USA

E-mail address: `spupyrev@gmail.com`