Drawing Permutations with Few Corners

Sergey Bereg[1], Alexander E. Holroyd[2], Lev Nachmanson[2], and Sergey Pupyrev[3,4,⋆]

[1] Department of Computer Science, University of Texas at Dallas, USA
[2] Microsoft Research, USA
[3] Department of Computer Science, University of Arizona, USA
[4] Institute of Mathematics and Computer Science, Ural Federal University, Russia

Abstract. A permutation may be represented by a collection of paths in the plane. We consider a natural class of such representations, which we call tangles, in which the paths consist of straight segments at 45 degree angles, and the permutation is decomposed into nearest-neighbour transpositions. We address the problem of minimizing the number of crossings together with the number of corners of the paths, focusing on classes of permutations in which both can be minimized simultaneously. We give algorithms for computing such tangles for several classes of permutations.

1 Introduction

What is a good way to visualize a permutation? In this paper we study drawings in which a permutation of interest is connected to the identity permutation via a sequence of intermediate permutations, with consecutive elements of the sequence differing by one or more non-overlapping nearest-neighbour swaps. The position of each permutation element through the sequence may then traced by a piecewise-linear path comprising segments that are vertical and $45°$ to the vertical. Our goal is to keep these paths as simple as possible and to avoid unnecessary crossings.

Such drawings have applications in various fields; for example, in channel routing for integrated circuit design [12]. Another application is the visualization of metro maps and transportation networks, where some lines (railway tracks or roads) might partially overlap [4]. A natural goal is to draw the lines along their common subpaths so that an individual line is easy to follow; minimizing the number of bends of a line and avoiding unnecessary crossings between lines are natural criteria for map readability; see Fig. 3(b) of [3]. Much recent research in the graph drawing community is devoted to edge bundling. In this setting, drawing the edges of a bundle with the minimum number of crossings and bends occurs as a subproblem [10].

Let S_n be the symmetric group of permutations $\pi = [\pi(1), \dots, \pi(n)]$ on $\{1, \dots, n\}$. The **identity permutation** is $[1, \dots, n]$, and the **swap** $\sigma(i)$ transforms a permutation π into $\pi \cdot \sigma(i)$ by exchanging its ith and $(i + 1)$th elements. Equivalently, $\sigma(i)$ is the transposition $(i, i + 1) \in S_n$, and $\cdot$ denotes composition. Two permutations a and b of S_n are **adjacent** if b can be obtained from a by swaps $\sigma(p_1), \sigma(p_2), \dots, \sigma(p_k)$ that are not overlapping, that is, such that $|p_i - p_j| \geq 2$ for $i \neq j$. A **tangle** is a finite sequence

⋆ Research supported in part by NSF grant DEB 1053573.

S. Wismath and A. Wolff (Eds.): GD 2013, LNCS 8242, pp. 484–495, 2013.

3	6	1	4	7	2	5
3	6	1	4	7	2	5
3	6	1	4	2	7	5
3	1	6	2	4	5	7
1	3	2	6	4	5	7
1	2	3	4	6	5	7
1	2	3	4	5	6	7
1	2	3	4	5	6	7

(a)

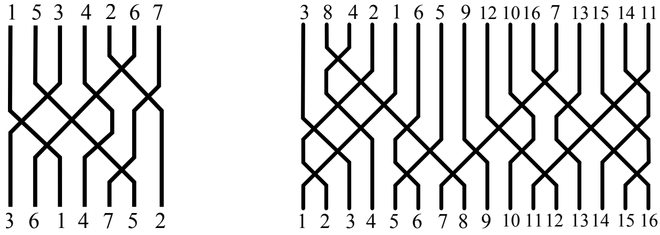

Fig. 1. (a) A tangle solving the permutation $[3, 6, 1, 4, 7, 2, 5]$. (b) A drawing of the tangle. (c) An example of a perfect tangle drawing.

of permutations in which each two consecutive permutations are adjacent. An example of a tangle is given in Fig. 1. The associated drawing is composed of polylines with vertices in $\mathbb{Z}^2$, whose segments can be vertical, or have slopes of $\pm 45°$ to the vertical. The polyline traced by element $i \in \{1, \dots, n\}$ is called **path** i. Note that by definition all path crossings occur at right angles. We say that a tangle T **solves** the permutation π (or simply T is a tangle for π) if the tangle starts from π and ends at the identity permutation.

We are interested in tangles with informative and aesthetically pleasing drawings. Our main criterion is to keep the paths straight by using only a few turns. A **corner** of path i is a point at which it changes its direction from one of the allowed directions (vertical, $+45°$, or $-45°$) to another. A change between $+45°$ and $-45°$ is called a **double corner**. We are interested in the total number of corners of a tangle, where corners are always counted with multiplicity (so a double corner contributes 2 to the total). By convention we require that paths start and end with vertical segments. In terms of the sequence of permutations this means repeating the first and the last permutations at least once each as in Fig. 1(a).

Another natural objective is to minimize path crossings. We call a tangle for π **simple** if it has the minimum number of crossings among all tangles for π. This is equivalent to the condition that no pair of paths cross each other more than once, and this minimum number equals the *inversion number* of π. A simple tangle has no double corner since that would entail an immediate double crossing of a pair of paths.

In general, minimizing corners and minimizing crossings are conflicting goals. For example, let $n = 4k$ and $k \geq 4$ and consider the permutation

$$\pi = [2k, 3, 2, 5, 4, \dots, 2k-1, 2k-2, 1, \quad 4k, 2k+3, 2k+2, \dots, 4k-1, 4k-2, 2k+1].$$

It is not difficult to check that the minimum number of corners in a tangle for π is $4n-8$, while the minimum among simple tangles is $5n-20$, which is strictly greater; see Fig. 2 for the case $k = 4$. Our focus in this article is on two special classes of permutations for which corners and crossings can be minimized simultaneously. The first is relatively straightforward, while the second turns out to be much more subtle.

One may ask the following interesting question. Is there an efficient algorithm for *finding a (simple) tangle with the minimum number of corners solving a given permutation*?

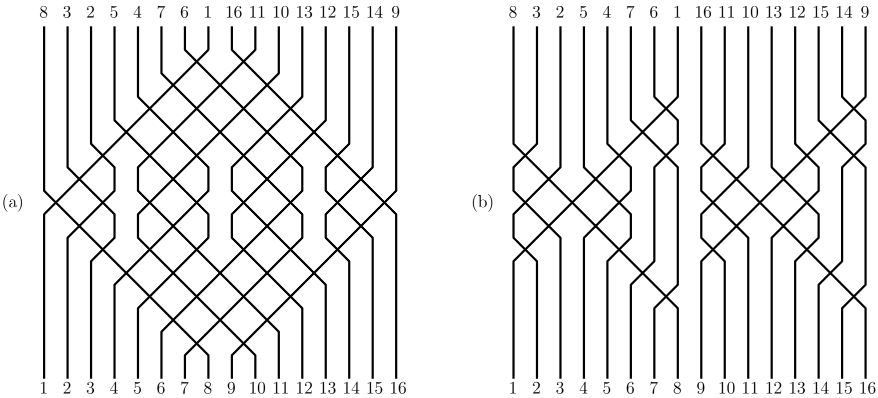

Fig. 2. (a) A tangle with 56 corners. (b) Every *simple* tangle for the same permutation has at least 60 corners.

We do not know whether there is a polynomial-time algorithm, either with or without the requirement of simplicity. Here we present polynomial-time exact algorithms for special classes of permutations.

Even the task of determining whether a given tangle has the minimum possible number of corners among tangles for its permutation does not appear to be straightforward in general (and likewise if we restrict to simple tangles). However, in certain cases, such minimality is indeed evident, and we focus on two such cases. Firstly, we call a tangle **direct** if each of its paths has at most 2 corners (equivalently, at most one non-vertical segment). Note that a direct tangle is simple. Furthermore, it clearly has the minimum number of corners among all tangles (simple or otherwise) for its permutation.

We can completely characterize permutations admitting direct tangles. We say that a permutation $\pi \in S_n$ **contains** a pattern $\mu \in S_k$ if there are integers $1 \leq i_1 < i_2 < \cdots < i_k \leq n$ such that for all $1 \leq r < s \leq k$ we have $\pi(i_r) < \pi(i_s)$ if and only if $\mu(r) < \mu(s)$; otherwise, π **avoids** the pattern (or π is μ-avoiding).

Theorem 1. *A permutation has a direct tangle if and only if it is* 321*-avoiding.*

Our proof yields a straightforward algorithm that constructs a direct tangle for a given 321-avoiding permutation.

Our second special class of tangles naturally extends the notion of a direct tangle, but turns out to have a much richer theory. A **segment** is a straight line segment of a path between two of its corners; it is an **L-segment** if it is oriented from north-east to south-west, and an **R-segment** if it is oriented from north-west to south-east. We call a tangle **perfect** if it is simple and each of its paths has at most one L-segment and at most one R-segment. Any perfect tangle has the minimum possible number of corners among all tangles solving its permutation, and indeed it has the minimum possible corners on path i for each $i = 1, \ldots, n$. To see this, note that if i has an L-segment in a perfect tangle for π then there must be an element $j > i$ with $\pi(i) > \pi(j)$, whose path crosses this L-segment. Hence, an L-segment must be present in any tangle for π. The same argument applies to R-segments. We call a permutation **perfect** if it has a perfect tangle.

Theorem 2. *There exists a polynomial-time algorithm that determines whether a given permutation is perfect and, if so, outputs a perfect tangle.*

A straightforward implementation of our algorithm takes $O(n^5)$ time, but we believe this can be reduced to $O(n^3)$, and possibly further. Our proof of Theorem 2 involves an explicit characterization of perfect permutations, but it is considerably more complicated than in the case of direct tangles. We introduce the notion of a *marking*, which is an assignment of symbols to the elements $1, \ldots, n$ indicating the directions in which their paths should be routed. We prove that a permutation is perfect if and only if it admits a marking satisfying a *balance* condition that equates numbers of elements in various categories. Finally, we show that the existence of such a marking can be decided by finding a maximum vertex-weighted matching in a certain graph with vertex set $1, \ldots, n$ constructed from the permutation.

The number of perfect permutations in S_n grows only exponentially with n (see Section 4), and is therefore $o(|S_n|)$. Nonetheless, perfect permutations are very common for small n: all permutations in S_6 are perfect, as are all but 16 in S_7, and over half in S_{13}.

Related Work. We are not aware of any other study on the number of corners in a tangle. To the best of our knowledge, the problem formulated here is new. Wang in [12] considered the same model of drawings in the field of VLSI design. However, [12] targets, in our terminology, the tangle height and the total length of the tangle paths. The heuristic suggested by Wang produces paths with many unnecessary corners.

The perfect tangle problem is related to the problem of drawing graphs in which every edge is represented by a polyline with few bends. In our setting, all the crossings occur at right angles, as in so-called RAC-drawings [6].

Decomposition of permutations into nearest-neighbour transpositions was considered in the context of permuting machines and pattern-restricted classes of permutations [1]. In our terminology, Albert et. al. [1] proved that it is possible to check in polynomial time whether for a given permutation there exists a tangle of length k (that is, consisting of k permutations), for a given k. Tangle diagrams appear in the drawings of sorting networks [8,2]. We also mention an interesting connection with change ringing (English-style church bell ringing), where similar visualizations are used [13].

2 Preliminaries

We always draw tangles oriented downwards with the sequence of permutations read from top to bottom as in Fig. 1(b). The following notation will be convenient. We write $\pi = [\ldots a \ldots b \ldots c \ldots]$ to mean that $\pi^{-1}(a) < \pi^{-1}(b) < \pi^{-1}(c)$, and $\pi = [\ldots ab \ldots]$ to mean that $\pi^{-1}(a) + 1 = \pi^{-1}(b)$. A pair of elements (a, b) is an **inversion** in a permutation $\pi \in S_n$ if $a > b$ and $\pi = [\ldots a \ldots b \ldots]$. The **inversion number** $\text{inv}(\pi) \in [0, \binom{n}{2}]$ is the number of inversions of π. The following useful lemma is straightforward to prove.

Lemma 1. *In a simple tangle for permutation π, a pair (i, j) is an inversion in π if and only if some R-segment of path i intersects some L-segment of path j.*

3 Direct Tangles

Here we prove Theorem 1. We need two properties of 321-avoiding permutations.

Lemma 2. *Suppose* π, π' *are permutations with* $\mathrm{inv}(\pi') = \mathrm{inv}(\pi) - 1$ *and* $\pi' = \pi \cdot \sigma(i)$ *for some swap* $\sigma(i)$*. If* π *is 321-avoiding then so is* π'*.*

Proof. Let us suppose that elements i, j, k form a 321-pattern in π'. Then (i, j) and (j, k) are inversions in π'. Inversions of π' are inversions of π, hence, elements i, j, k form a 321-pattern in π. □

Lemma 3. *In a simple tangle solving a 321-avoiding permutation, no path has both an L-segment and an R-segment.*

Proof. Consider a simple tangle solving a 321-avoiding permutation π. Suppose path j crosses path i during j's R-segment and crosses path k during j's L-segment. By Lemma 1 we have $\pi = [\dots k \dots j \dots i \dots]$ while $i < j < k$, giving a 321-pattern, which is a contradiction. □

We say that a permutation $\pi \in S_n$ has a **split** at location k if $\pi(1), \dots, \pi(k) \in \{1, \dots, k\}$, or equivalently if $\pi(k+1), \dots, \pi(n) \in \{k+1, \dots, n\}$.

Theorem 1. *A permutation has a direct tangle if and only if it is* 321*-avoiding.*

Proof. To prove the "only if" part, suppose that tangle T solves a permutation π containing a 321-pattern. Then there are $i < j < k$ with $\pi = [\dots k \dots j \dots i \dots]$. Hence by Lemma 1, j has an L-segment and an R-segment, so T is not direct.

We prove the "if" part by induction on the inversion number of the permutation. If $\mathrm{inv}(\pi) = 0$ then π is the identity permutation, which clearly has a direct tangle. This gives us the basis of induction.

Now suppose that π is 321-avoiding and not the identity permutation, and that every 321-avoiding permutation (of every size) with inversion number less than $\mathrm{inv}(\pi)$ has a direct tangle. There exists s such that $\pi(s) > \pi(s+1)$; fix one such. Note that $(\pi(s), \pi(s+1))$ is an inversion of π; hence, the permutation $\pi' := \pi \cdot \sigma(s)$ has $\mathrm{inv}(\pi') = \mathrm{inv}(\pi) - 1$, and is also 321-avoiding by Lemma 2. By the induction hypothesis, let T' be a direct tangle solving π'.

Perform a swap x in position s exchanging elements $\pi(s)$ and $\pi(s+1)$, and draw it as a cross on the plane with coordinates (s, h), where $h \in \mathbb{Z}$ is the height (y-coordinate) of the cross (chosen arbitrarily). We assume that the position axis increases from left to right and the height axis increases from bottom to top. Then draw the tangle T' below the cross. This gives a tangle solving π, which is certainly simple. We show that the heights of swaps may be adjusted to make the new tangle direct. To achieve this, the L-segment and R-segment comprising the swap x must either extend existing segments in T', or must connect to vertical paths having no corners in T'. Consider two cases.

Case 1: Suppose that π' has a split at s. Then T' consists of a tangle T_1 for the permutation $[\pi'(1), \dots, \pi'(s)]$ together with another tangle T_2 for $[\pi'(s+1), \dots, \pi'(n)]$; see Fig. 3. Starting with T_1 drawn below x, simultaneously shift all the swaps of T_1 upward until one of them touches x; in other words, until T_1's first swap in position

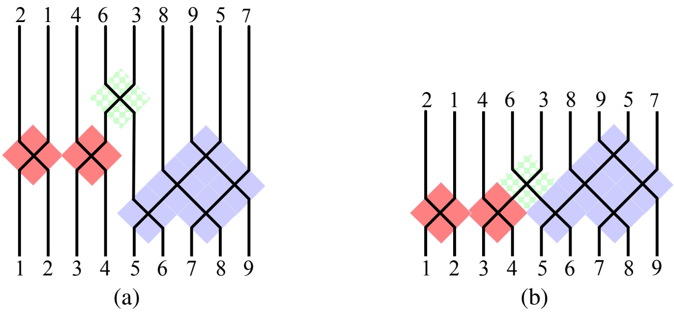

Fig. 3. Shifting two sub-tangles (T_1 is red, T_2 is blue) upward to touch the initial swap x (green) in position $s = 4$.

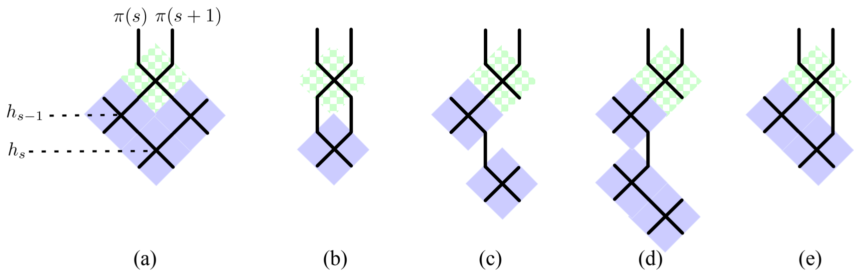

Fig. 4. (a) The tangle T' (blue) touches the swap x (green) on both sides. (b)–(e) Various impossible configurations for the proof.

$s - 1$ occurs at height $h - 1$. Or, if T_1 has no swap in position $s - 1$, no shifting is necessary. Similarly shift T_2 upward until it touches x from the right side. This results in a direct tangle.

Case 2: Suppose that π' has no split at s. Let T' be any direct tangle for π', and again shift it upward until it touches x, resulting in a tangle T for π. Write h_j for the height of the topmost swap in position j in T', or let $h_j = -\infty$ if there is none. We claim that $h_{s-1} = h_{s+1} = h_s + 1 > -\infty$, which implies in particular that $1 < s < s + 1 < n$. Thus T' has swaps in the positions immediately left and right of x, both of which touch x simultaneously in the shifting procedure as in Fig. 4(a), giving that T is direct as required. To prove the claim, first note that $h_s > -\infty$ since π' has no split at s. Therefore $\max\{h_{s-1}, h_{s+1}\} > h_s$, otherwise T would not be simple, as in Fig. 4(b). Thus, without loss of generality suppose that $h_{s-1} > h_s$ and $h_{s-1} \geq h_{s+1}$. Then $h_s = h_{s-1} - 1$, otherwise some path would have more than 2 corners in T', specifically, the path of the element that is in position s after h_{s-1}; see Fig. 4(c) or (d). Now suppose for a contradiction that $h_{s+1} < h_{s-1}$, which includes the possibility that $h_{s+1} = -\infty$, perhaps because $s + 1 = n$. Then in the new tangle T, path $\pi(s)$ contains both an L-segment and an R-segment as in Fig. 4(e), which contradicts Lemma 3. □

The proof of Theorem 1 yields an algorithm that returns a direct tangle for $\pi \in S_n$ if one exists, and otherwise stops. The algorithm can be implemented so as to run in $O(n^2)$ time. With a suitable choice of output format, this can be improved to $O(n)$.

4 Perfect Tangles

In this section we give our characterization of perfect permutations. Given a permutation $\pi \in S_n$, we introduce the following classification scheme of elements $i \in \{1, \ldots, n\}$. The scheme reflects the possible forms of paths in a perfect tangle, although the definitions themselves are purely in terms of the permutation. We call i a **right** element if it appears in some inversion of the form (i, j), and a **left** element if it appears in some inversion (j, i). We call i **left-straight** if it is left but not right, **right-straight** if it is right but not left, and a **switchback** if it is both left and right.

In order to build a perfect tangle we use a notion of marking. A **marking** M is a function from the set $\{1, \ldots, n\}$ to strings of letters L and R. For any tangle T, we associate a corresponding marking M as follows. We trace the path i from top to bottom; as we meet an L-segment (resp. R-segment), we append an L (resp. R) to $M(i)$. Vertical segments are ignored for this purpose; hence, a vertical path with no corners is marked by an empty sequence $\emptyset$. For example, $M(3) = R$ and $M(13) = LR$ in Fig. 1(c). A marking corresponding to a perfect tangle takes only values $\emptyset$, L, R, LR, and RL. We write $M(i) = R\ldots$ to indicate that the string $M(i)$ starts with R.

Given a permutation π and a marking M, there does not necessarily exist a corresponding tangle. However, we will obtain a necessary and sufficient condition on π and M for the existence of a corresponding perfect tangle. Our strategy for proving Theorem 2 will be to find a marking satisfying this condition, and then to find a corresponding perfect tangle. We say that a marking M is a **marking for** a permutation $\pi \in S_n$ if (i) $M(i) = L$ (respectively $M(i) = R$) for all left-straight (right-straight) elements i, (ii) $M(i) \in \{LR, RL\}$ for all switchbacks, and (iii) $M(i) = \emptyset$ otherwise.

To state the necessary and sufficient condition mentioned above, we need some definitions. A quadruple (a, b, c, d) is a **rec** in permutation π if $\pi = [\ldots a \ldots b \ldots c \ldots d \ldots]$ and $\min\{a, b\} > \max\{c, d\}$. In a perfect tangle, the paths comprising a rec form a rectangle; see Fig. 5 ("rec" is an abbreviation for rectangle). Let M be a marking for $\pi \in S_n$, and let ρ be a rec (a, b, c, d) in π. We call e a **left switchback** of ρ if (i) $M(e) = RL$, (ii) $\pi = [\ldots a \ldots e \ldots b \ldots]$, and (iii) $c < e < d$ or $d < e < c$. Symmetrically, we call e a **right switchback** of ρ if $M(e) = LR$, and $\pi = [\ldots c \ldots e \ldots d \ldots]$, and $a < e < b$ or $b < e < a$. A rec (a, b, c, d) is **regular** if $a < b$ and $c < d$, otherwise it is **irregular**. A rec is called **balanced** under M if the number of its left switchbacks is equal to the number of its right switchbacks; a rec is **empty** if it has no switchbacks.

Here is our key definition. A marking M for a permutation π is called **balanced** if every regular rec of π is balanced and every irregular rec is empty under M.

Theorem 3. *A permutation is perfect if and only if it admits a balanced marking.*

The proof of Theorem 3 is technical, see full version for the complete proof [5].

Any permutation containing the pattern $[7324651]$ (for example) is not perfect since 4 must be a switchback of one of the irregular recs (7321) and (7651). It follows by [9,7]

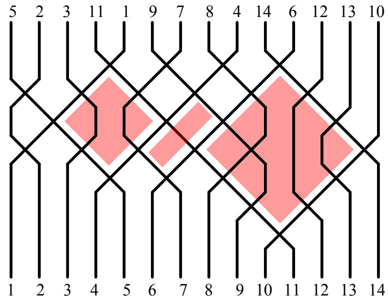

Fig. 5. A permutation with a balanced marking. Some of the recs of the permutation are: $\rho_1 = (5, 11, 1, 4)$, $\rho_2 = (9, 7, 4, 6)$, $\rho_3 = (11, 14, 6, 10)$; ρ_1 and ρ_3 are regular, while ρ_2 is irregular. Left switchbacks of rec ρ_3 are 8 and 9, right switchbacks are 12 and 13. The empty irregular rec ρ_2 has neither left nor right switchbacks.

that the number of perfect permutations in S_n is at most C^n for some constant $C > 1$. Since direct tangles are perfect, it also follows from Theorem 1 that the number is at least c^n for some constant $c > 1$.

We note that Theorem 3 already yields an algorithm for determining whether a permutation is perfect in $\widetilde{O}(2^n)$ time[1] by checking all markings. In Section 5 we improve this to polynomial time.

5 Recognizing Perfect Permutations

We provide an algorithm for recognizing perfect permutations. The algorithm finds a balanced marking for a permutation, or reports that such a marking does not exist. We start with a useful lemma.

Lemma 4. *Fix a permutation. For each right (resp., left) element a there is a left-straight (right-straight) b such that the pair (a, b) (resp., (b, a)) is an inversion.*

Proof. We prove the case when a is right, the other case being symmetrical. Consider the minimal b such that (a, b) is an inversion. By definition, b is left. Suppose that it is also a right element, that is, (b, c) is an inversion for some $c < b$. It is easy to see that (a, c) is an inversion too, which contradicts to the minimality of b. □

Recall that a marking is balanced only if (in particular) every regular rec of the permutation is balanced under the marking. We show that this is guaranteed even by balancing of recs of a restricted kind. We call a rec (a, b, c, d) of a permutation π **straight** if a, b, c, and d are straight elements of π. A marking is called **s-balanced** if every straight rec is balanced and every irregular rec is empty under the marking.

Lemma 5. *Let M be a marking of a permutation π. Then M is balanced if and only if it is s-balanced.*

[1] $\widetilde{O}$ hides a polynomial factor.

Proof. The "if" direction is immediate, so we turn to the converse. Let M be an s-balanced marking and $\rho = (a, b, c, d)$ be a regular rec of π. We need to prove that ρ is balanced under M. If ρ is straight then ρ is balanced by definition. Let us suppose that ρ is not straight. Then some $u \in \{a, b, c, d\}$ is not a straight element. Our goal is to show that it is possible to find a new rec ρ' in which u is replaced with a straight element so that the sets of left and right switchbacks of ρ and ρ' coincide. By symmetry, we need only consider the cases $u = a$ and $u = b$.

Case $u = a$**:** Let us suppose that a is not straight. By Lemma 4, there exists a right straight e such that (e, a) is an inversion. Let us denote $\rho' = (e, b, c, d)$ and show that ρ' has the same switchbacks as ρ. Let k be a left switchback of ρ; then $M(k) = RL$, and $\pi = [\dots e \dots a \dots k \dots b \dots]$, and $c < k < d$. By definition k is a left switchback of ρ'. Let k be a left switchback of ρ'. If $\pi = [\dots e \dots k \dots a \dots b \dots]$ then the irregular rec (e, a, c, d) has a left switchback, which is impossible. Therefore, $\pi = [\dots e \dots a \dots k \dots b \dots]$ and k is a left-switchback of ρ.

Let us suppose that k is a right switchback of ρ, so $a < k < b$. If $k < e$ then k is a right switchback of the irregular (e, a, c, d); hence, $e < k < b$ and k is a right switchback of ρ'. On the other hand, if k is a right switchback of ρ' then $a < e < k < b$, which means that k is a right switchback of ρ.

Case $u = b$**:** Let us suppose that b is not straight. By Lemma 4, there exists a right straight e such that (e, b) is an inversion. Let us denote $\rho' = (a, e, c, d)$ and show that ρ' has the same switchbacks as ρ. Let k be a left switchback of ρ. We have $\pi = [\dots a \dots k \dots b \dots]$. Since k is not a left switchback of the irregular rec (e, b, c, d), we have $\pi = [\dots a \dots k \dots e \dots]$. Therefore, k is a left switchback of ρ'.

Let k be a right switchback of ρ. Then $a < k < b < e$, proving that k is a right switchback of ρ'. Let k be a right switchback of ρ'. If $b < k$ then k is a right switchback of (e, b, c, d), which is impossible. Then $k < b$ and k is a right switchback of ρ. □

We can restrict the set of recs guaranteeing the balancing of a permutation even further. We call a pair a, b of elements **right** (resp. **left**) **minimal** if a and b are right (left) straight elements of π, and $a < b$, and there is no right (left) straight element c such that $\pi = [\dots a \dots c \dots b \dots]$. We call rec $\rho = (a, b, c, d)$ **minimal** in π if a, b is a right minimal pair and c, d is a left minimal pair; see Fig. 6(a). We call a marking for a permutation **ms-balanced** if every minimal regular rec is balanced and every irregular rec is empty under the marking.

Lemma 6. *Let M be a marking of a permutation π. Then M is s-balanced if and only if it is ms-balanced.*

Before giving the proof, we introduce some further notation. Let $\rho = (a, b, c, d)$ be an arbitrary, possibly irregular, rec in π. Let us denote by ρ_ℓ (resp. ρ_r) the set of switchbacks i that can under some marking be left (resp., right) switchbacks of ρ. Formally, $i \in \rho_\ell$ if and only if $\pi = [\dots a \dots i \dots b \dots c \dots d \dots]$ and either $c < i < d$ or $d < i < c$. (And ρ_r is defined symmetrically.) For a rec ρ and marking M let ρ_ℓ^M (ρ_r^M) be the set of left (respectively, right) switchbacks of ρ under M. Of course, $\rho_\ell^M \subseteq \rho_\ell$ and $\rho_r^M \subseteq \rho_r$. It is easy to see from the definition that for two different minimal recs ρ and ρ' we have $\rho_\ell \cap \rho'_\ell = \emptyset$ and $\rho_r \cap \rho'_r = \emptyset$.

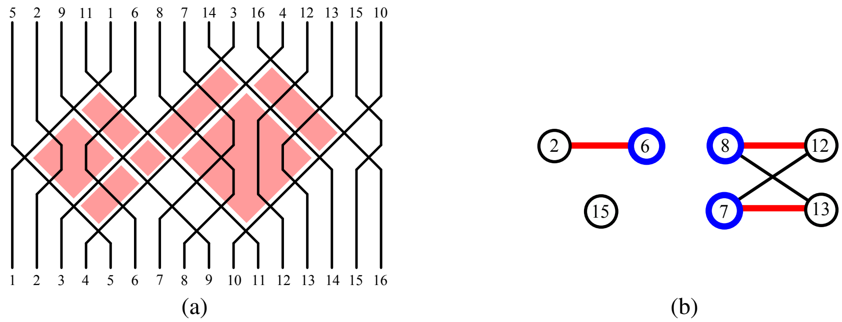

Fig. 6. (a) A perfect tangle for a permutation with 7 minimal straight recs (shown red). (b) The graph constructed in *Step 3* of our algorithm. Here, $\mathfrak{I}_\ell = \emptyset$, $\mathfrak{I}_r = \{6\}$, $\mathfrak{R}_\ell = \{6, 7, 8, 12, 13, 15\}$, and $\mathfrak{R}_r = \{2, 7, 8\}$. The vertices of the set $F = \{6, 7, 8\}$ are shown blue. The red edges are the computed maximum matching.

Proof (Lemma 6). It suffices to prove that if M is ms-balanced then it is s-balanced. Consider a straight rec $\rho = (a, b, c, d)$. Let $a = r_1, \dots, r_p = b$ be a sequence of right straights in which each consecutive pair r_i, r_{i+1} is right minimal. Define left straights $c = \ell_1, \dots, \ell_q = d$ similarly. Let D be the set of all recs of the form $(r_i, r_{i+1}, \ell_j, \ell_{j+1})$ for $1 \le i < p$ and $1 \le j < q$. Notice that all recs of D are minimal. By definition of rec switchbacks, we have $\rho_\ell^M = \bigcup_{u \in D} u_\ell^M$ and $\rho_r^M = \bigcup_{u \in D} u_r^M$. Since every rec $u \in D$ is balanced and for every pair $u, v \in D$ of different recs $u_\ell^M \cap v_\ell^M = u_r^M \cap v_r^M = \emptyset$, we have $|\rho_\ell^M| = |\rho_r^M|$; that is, ρ is balanced under M. $\square$

Let us show how to construct an ms-balanced marking. For a permutation π, let $\mathfrak{I}_\ell = \bigcup\{\rho_\ell : \rho \text{ is an irregular rec in } \pi\}$ and $\mathfrak{R}_\ell = \bigcup\{\rho_\ell : \rho \text{ is a regular rec in } \pi\}$, and define $\mathfrak{I}_r, \mathfrak{R}_r$ similarly. Our algorithm is based on finding a maximum vertex-weighted matching, which can be done in polynomial time [11].

The algorithm inputs a permutation π and computes an ms-balanced marking M for π or determines that such a marking does not exist. Initially, $M(i)$ is undefined for every $i \in \{1, \dots, n\}$. The algorithm has the following steps.

***Step 1*:** For every element $1 \le i \le n$ that is neither left nor right, set $M(i) = \emptyset$. For every left straight i set $M(i) = L$. For every right straight i set $M(i) = R$.

***Step 2*:** If $\mathfrak{I}_\ell \cap \mathfrak{I}_r \ne \emptyset$ then report that π is not perfect and stop. Otherwise, for every switchback $i \in \mathfrak{I}_\ell$ set $M(i) = LR$; for every switchback $i \in \mathfrak{I}_r$ set $M(i) = RL$.

***Step 3.1*:** Build a directed graph $G = (V, E)$ with $V = \mathfrak{R}_\ell \cup \mathfrak{R}_r$ and $E = \bigcup\{(\rho_\ell \setminus \mathfrak{I}_\ell) \times (\rho_r \setminus \mathfrak{I}_r) : \rho \text{ is a minimal rec in } \pi\}$.

***Step 3.2*:** Create a set $F \leftarrow (\mathfrak{R}_\ell \cap \mathfrak{R}_r) \cup (\mathfrak{I}_\ell \cap \mathfrak{R}_r) \cup (\mathfrak{I}_r \cap \mathfrak{R}_\ell)$. Create weights w for vertices of G: if $i \in F$ then set $w(i) = 1$, otherwise set $w(i) = 0$.

***Step 4*:** Compute a maximum vertex-weighted matching U on G (viewed as an *unoriented* graph, ignoring the directions of edges) using weights w. If the total weight of U is less than $|F|$ then report that π is not perfect and stop.

Step 5.1: Assign marking based on the matching: for every edge $(i, j) \in U$ set $M(i) = RL$ provided $M(i)$ has not already been assigned, and $M(j) = LR$ provided $M(j)$ has not already been assigned.

Step 5.2: For every switchback $1 \leq i \leq n$ with still undefined marking, if $i \in \mathfrak{R}_\ell$ then set $M(i) = LR$, if $i \in \mathfrak{R}_r$ then $M(i) = RL$, otherwise choose $M(i)$ to be LR or RL arbitrarily. Note that any $i \in \mathfrak{R}_\ell \cap \mathfrak{R}_r$ was already assigned because of *Steps 3.2 and 4*.

Let us prove the correctness of the algorithm.

Lemma 7. *If the algorithm produces a marking then the marking is ms-balanced.*

Proof. Let M be a marking produced by the algorithm for a permutation π. It is easy to see that $M(i)$ is defined for all $1 \leq i \leq n$ (in *Step 1* for straights and in *Step 2* and *Step 5* for switchbacks). By construction, M is a marking for π.

Let us show that M is ms-balanced. Consider an irregular rec ρ of π, and suppose that $i \in \rho_\ell$. Since $\rho_\ell \subseteq \mathfrak{I}_\ell$, in *Step 2* we assign $M(i) = LR$, that is, $i \notin \rho_\ell^M$. Therefore, ρ does not have left switchbacks under M. Similarly, ρ does not have right switchbacks under M. Therefore, ρ is empty.

Consider a regular minimal straight rec ρ in π. Suppose that $i \in \rho_\ell^M$. Then $M(i) = RL$ and $i \in \rho_\ell \subseteq \mathfrak{R}_\ell$. If $i \in \mathfrak{I}_r$ then $i \in \mathfrak{R}_\ell \cap \mathfrak{I}_r \subseteq F$; hence i is incident to an edge in U. Since no directed edge of the form (k, i) is included in G in *Step 3.1*, there exists $(i, k) \in U$ for some k. On the other hand, if $i \notin \mathfrak{I}_r$ then string RL was not assigned to $M(i)$ in *Step 5.2*, nor in *Step 2*. Thus, it was assigned in *Step 5.1*, and again $(i, k) \in U$ for some k. By definition of E we have $k \in \rho_r$, because k cannot appear in ρ'_r for any other minimal $\rho' \neq \rho$. The algorithm sets $M(k) = LR$ at *Step 5.1*; it could not have previously set $M(k) = LR$ at *Step 2* because $k \notin \mathfrak{I}_r$ by the definition of E. Thus $k \in \rho_r^M$.

By symmetry, an identical argument to the above shows that if $k \in \rho_r^M$ then $i \in \rho_\ell^M$ for some i satisfying $(i, k) \in U$. Since U is a matching, we thus have a bijection between elements of ρ_ℓ^M and ρ_r^M. Therefore, ρ is balanced under M. □

Lemma 8. *Let π be a perfect permutation. The algorithm produces a marking for π.*

Proof. Since π is perfect, there is a balanced marking M for π. Since M is balanced, all irregular recs are empty under M; hence, the algorithm does not stop in *Step 2*. To prove the claim, we will create a matching in the graph G with total weight $|F|$.

Let ρ be a minimal rec in π. Since ρ is balanced under M, we have $|\rho_\ell^M| = |\rho_r^M|$. Hence, let W_ρ be an arbitrary matching connecting vertices of $|\rho_\ell^M|$ with vertices of $|\rho_r^M|$. Of course, $|W_\rho| = |\rho_\ell^M|$. Let $W = \bigcup\{W_\rho : \rho \text{ is a minimal rec in } \pi\}$. We show that every element of set F is incident to an edge of W.

Suppose $i \in \mathfrak{R}_\ell \cap \mathfrak{R}_r$. Since i is a switchback in π, we have $M(i) = RL$ or $M(i) = LR$. In the first case $i \in \rho_\ell^M$ and in the second case $i \in \rho_r^M$ for some minimal rec ρ. Then i is incident to an edge from W_ρ.

Suppose $i \in F \setminus \{\mathfrak{R}_\ell \cap \mathfrak{R}_r\}$. Without loss of generality, let $i \in \mathfrak{I}_\ell \cap \mathfrak{R}_r$. Since M is balanced, every irregular rec has no switchbacks and hence $M(i) = LR$. Thus, $i \in \rho_r^M$ for some minimal rec ρ, and i is incident to an edge of W_ρ.

Therefore, every vertex of F is incident to an edge of the matching W, which means that the total weight of W is $|F|$. □

Theorem 2 follows directly from Lemmas 7 and 8 and Theorem 3. A straightforward implementation of the algorithm finding a perfect tangle takes $O(n^5)$ time.

6 Conclusion

In this paper we gave algorithms for producing optimal tangles in the special cases of direct and perfect tangles, and for recognizing permutations for which this is possible. Many questions remain open. What is the complexity of determining the tangle with minimum corners for a given permutation? What is the complexity if the tangle is required to be simple? What is the asymptotic behavior of the maximum over permutations $\pi \in S_n$ of the minimum number of corners among simple tangles solving π?

Acknowledgments: We thank Omer Angel, Franz Brandenburg, David Eppstein, Martin Fink, Michael Kaufmann, Peter Winkler, and Alexander Wolff for fruitful discussions about variants of the problem.

References

1. Albert, M.H., Aldred, R.E.L., Atkinson, M., van Ditmarsch, H.P., Handley, C.C., Holton, D.A., McCaughan, D.J.: Compositions of pattern restricted sets of permutations. Australian J. Combinatorics 37, 43–56 (2007)
2. Angel, O., Holroyd, A.E., Romik, D., Virag, B.: Random sorting networks. Advances in Mathematics 215(2), 839–868 (2007)
3. Argyriou, E.N., Bekos, M.A., Kaufmann, M., Symvonis, A.: On metro-line crossing minimization. Journal of Graph Algorithms and Applications 14(1), 75–96 (2010)
4. Benkert, M., Nöllenburg, M., Uno, T., Wolff, A.: Minimizing intra-edge crossings in wiring diagrams and public transportation maps. In: Kaufmann, M., Wagner, D. (eds.) GD 2006. LNCS, vol. 4372, pp. 270–281. Springer, Heidelberg (2007)
5. Bereg, S., Holroyd, A.E., Nachmanson, L., Pupyrev, S.: Drawing permutations with few corners. ArXiv e-print abs/1306.4048 (2013)
6. Didimo, W., Eades, P., Liotta, G.: Drawing graphs with right angle crossings. Theoretical Computer Science 412(39), 5156–5166 (2011)
7. Klazar, M.: The Füredi-Hajnal conjecture implies the Stanley-Wilf conjecture. In: Krob, D., Mikhalev, A., Mikhalev, A. (eds.) Formal Power Series and Algebraic Combinatorics, pp. 250–255. Springer, Heidelberg (2000)
8. Knuth, D.: The art of computer programming. Addison-Wesley (1973)
9. Marcus, A., Tardos, G.: Excluded permutation matrices and the Stanley-Wilf conjecture. Journal of Combinatorial Theory, Series A 107(1), 153–160 (2004)
10. Pupyrev, S., Nachmanson, L., Bereg, S., Holroyd, A.E.: Edge routing with ordered bundles. In: van Kreveld, M.J., Speckmann, B. (eds.) GD 2011. LNCS, vol. 7034, pp. 136–147. Springer, Heidelberg (2012)
11. Spencer, T.H., Mayr, E.W.: Node weighted matching. In: Paredaens, J. (ed.) ICALP 1984. LNCS, vol. 172, pp. 454–464. Springer, Heidelberg (1984)
12. Wang, D.C.: Novel routing schemes for IC layout, part I: Two-layer channel routing. In: 28th ACM/IEEE Design Automation Conference, pp. 49–53 (1991)
13. White, A.T.: Ringing the changes. Mathematical Proceedings of the Cambridge Philosophical Society 94, 203–215 (1983)