Complexity Results on Untangling Planar Rectilinear Red-Blue Matchings

Arun Kumar Das*[1], Sandip Das*[2], Guilherme D. da Fonseca†[3], Yan Gerard†[4], and Bastien Rivier†[5]

1 **Indian Statistical Institute, Kolkata, India**
arund426@gmail.com
2 **Indian Statistical Institute, Kolkata, India**
sandipdas@isical.ac.in
3 **Aix-Marseille Université and LIS, France**
guilherme.fonseca@lis-lab.fr
4 **Université Clermont Auvergne and LIMOS, France**
yan.gerard@uca.fr
5 **Université Clermont Auvergne and LIMOS, France**
bastien.rivier@uca.fr

Abstract

Given a rectilinear matching between n red points and n blue points in the plane, we consider the problem of obtaining a crossing-free matching through flip operations that replace two crossing segments by two non-crossing ones. We first show that (i) it is NP-hard to α-approximate the shortest flip sequence, for any constant α. Second, we show that when the red points are colinear, (ii) given a matching, a flip sequence of length at most $\binom{n}{2}$ always exists, and (iii) the number of flips in any sequence never exceeds $\binom{n}{2}\frac{n+4}{6}$. Finally, we present (iv) a lower bounding flip sequence with roughly $1.5\binom{n}{2}$ flips, which disproves the conjecture that $\binom{n}{2}$, reached in the convex case, is the maximum. The last three results, based on novel analyses, improve the constants of state-of-the-art bounds.

Related Version arxiv.org/abs/2202.11857

1 Introduction

We consider the problem of untangling a planar rectilinear red-blue matching. We are given a set of $2n$ points in the plane, partitioned into a set R of n red points, and a set B of n blue points, in general position (no three colinear points, unless they have the same color).

A *configuration* is a set of n line segments where each point of R is matched to exactly one point of B, i.e. a perfect rectilinear red-blue matching. A flip is a combinatorial operation changing a configuration into another [7, 16]. In our case, a *flip* replaces two crossing segments by two non-crossing ones (Figure 1).

The *reconfiguration graph* of R, B is the directed simple graph whose vertices $\mathcal{V}$ are the configurations, and such that there is a directed edge from a configuration M_1 to another one M_2 whenever a flip transforms M_1 into M_2. Note that the reconfiguration graph is acyclic [6]. Let $\mathcal{S} \subseteq \mathcal{V}$ be the set of sinks, which corresponds to the crossing-free configurations. Given two configurations $u, v \in \mathcal{V}$, let $\mathcal{P}(u, v)$ be the set of directed paths from u to v. Given a path P, let the *length* of P, denoted $|P|$, be the number of edges in P. The *distance* from u

* This work is partially supported by the IFCAM project "Applications of graph homomorphisms" (MA/IFCAM/18/39).
† This work is supported by the French ANR PRC grant ADDS (ANR-19-CE48-0005).

38th European Workshop on Computational Geometry, Perugia, Italy, March 14–16, 2022.
This is an extended abstract of a presentation given at EuroCG'22. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

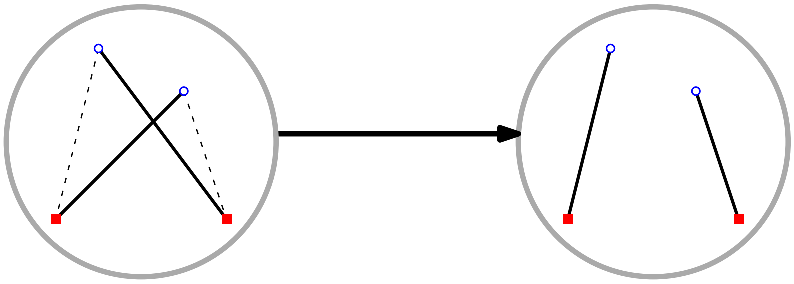

Figure 1 A flip. Red points are represented by solid squares and blue points by hollow circles.

to v, denoted $d(u, v)$, is the minimum path length from u to v. The *distance* from u to $\mathcal{S}$, $d(u, \mathcal{S})$, also abbreviated as $d(u)$, is the minimum path length from u to a configuration in $\mathcal{S}$. We are interested in two parameters of this reconfiguration graph:

$$\mathbf{d}(R, B) = \max_{u \in \mathcal{V}} \min_{v \in \mathcal{S}} \min_{P \in \mathcal{P}(u,v)} |P| \quad \text{and} \quad \mathbf{D}(R, B) = \max_{u \in \mathcal{V}} \max_{v \in \mathcal{S}} \max_{P \in \mathcal{P}(u,v)} |P| \,.$$

This leads to the definitions of $\mathbf{d}(n)$ and $\mathbf{D}(n)$ respectively as the maximum of $\mathbf{d}(R, B)$ and $\mathbf{D}(R, B)$ with $|R| = |B| = n$. An *untangle sequence* is a path in the reconfiguration graph ending in $\mathcal{S}$. Intuitively, $\mathbf{d}$ corresponds to the minimum length of an untangle sequence in the worst case, while $\mathbf{D}$ corresponds to the longest untangle sequence.

We also consider a more specific version of the problem where the red points are colinear [4], say, on the x-axis. As the flips on each half-plane defined by the x-axis are independent, we additionally suppose all blue points to lie on the upper half-plane without loss of generality. The matchings in this case are called *red-on-a-line* matchings.

Related work. The parameters $\mathbf{d}$, $\mathbf{D}$ have been studied in several different contexts with similar definitions of a flip, but considering other configurations.

In 1981, an $O(n^3)$ upper bound on $\mathbf{D}(n)$ was stated in the context of optimizing a TSP tour [23] (the configurations are polygons). This upper bound should be compared to the exponential lower bound on $\mathbf{D}(n)$ when the flips are not restricted to crossing segments, as long as they decrease the Euclidean length of the tour [10]. The convex case (i.e. the case where the points are in convex position) has been studied in [20, 25].

In the non-bipartite version of the rectilinear perfect matching problem, there are two possible pairs of segments to replace a crossing pair. This additional choice yields an $n^2/2$ upper bound on $\mathbf{d}(n)$ [6].

It is also possible to relax the flip definition to all operations that replace two segments by two others with the same four endpoints, whether they cross or not, and generalize the configurations to multigraphs with the same degree sequence [12, 13, 16]. In this context, finding the shortest path from a given configuration to another in the reconfiguration graph is NP-hard, yet 1.5-approximable [2, 3, 11, 24]. If we additionally require the configurations to be connected graphs, the same problem is NP-hard and 2.5-approximable [8].

Reconfiguration problems in the context of triangulations are widely studied [19]. A flip consists of removing one edge and adding another one while preserving a triangulation. It is know that $\Theta(n^2)$ flips are sufficient and sometimes necessary to obtain a Delaunay triangulation [14, 17]. Determining the flip distance between two triangulations of a point set [18, 21] and between two triangulations of a simple polygon [1] are both NP-hard.

Considering perfect matchings of an arbitrary graph (instead of the complete bipartite graph on R, B), a flip amounts to exchanging the edges in an alternating cycle of length

Table 1 Lower and upper bounds on $\mathbf{d}(n)$ and $\mathbf{D}(n)$ for red-blue matchings.

	$\mathbf{d}(n)$ bounds		$\mathbf{D}(n)$ bounds	
	lower	upper	lower	upper
general	$1.4n^{(a)}$, Thm. 5.2	$\binom{n}{2}(n-1)$, [6, 23]	$\frac{3}{2}\binom{n}{2} - \frac{n}{4}^{(b)}$, Thm. 5.1	$\binom{n}{2}(n-1)$, [6, 23]
convex	$1.4n^{(a)}$, Thm. 5.2	$2n-2$, [4]	$\binom{n}{2}$, [6]	$\binom{n}{2}$, [4]
red-on-a-line	$n-1$, [6]	$\binom{n}{2}$, Thm. 3.1	$\frac{3}{2}\binom{n}{2} - \frac{n}{4}^{(b)}$, Thm. 5.1	$\binom{n}{2}\frac{n+4}{6}$, Thm. 4.1

(a) For n multiple of 20.
(b) For even n.

four. It is then PSPACE-complete to decide whether there exists a path from a configuration to another [5]. There is, actually, a wide variety of reconfiguration contexts derived from NP-complete problems where this same accessibility problem is PSPACE-complete [15]. Many other reconfiguration problems are presented in [22].

Getting back to our context of rectilinear red-blue matchings, the values of $\mathbf{d}$ and $\mathbf{D}$ have been determined almost exactly in the convex case (see Table 1). Notice that the $n-1$ lower bound on $\mathbf{d}(n)$ carries to both the general and red-on-a-line cases [6]. It is notable that the upper bound on $\mathbf{D}(n)$ is also the best known bound on $\mathbf{d}(n)$ and has not been improved since 1981 [23].

Contributions. We show in Section 2 that it is NP-hard to α-approximate the shortest untangle sequence starting at a given matching, for any fixed $\alpha \geq 1$.

The following results are summarized in Table 1. An improved lower bound on $d(n)$ in the convex case is presented in Section 5.2. The remainder of the paper considers the red-on-a-line case. In Section 3, we slightly improve the former $\binom{n+1}{2}$ upper bound on $\mathbf{d}(n)$ [4], using a simpler algorithm and a novel analysis. In Section 4, we asymptotically divide by 6 the historical $\binom{n}{2}(n-1)$ upper bound on $\mathbf{D}(n)$ [6, 23], using a different potential argument.

In Section 5.1, we present a counter-example to the intuitive conjecture that the longest untangle sequence is attained in the convex case (where the number of crossings is maximal). We take advantage of points that are not in convex position to increase the lower bound by a factor of $\frac{3}{2}$. This red-on-a-line lower bound on $\mathbf{d}(n)$ carries over to the general case (and even to the case of general perfect matchings without color distinction among the points). The weaker conjecture that $\mathbf{D}(n)$ is quadratic [6] still holds, though.

2 NP-Hardness

In this section, we sketch the reduction of a known NP-complete problem, called rectilinear planar monotone (*RPM*) 3-SAT [9], to the following problem. The full proof is presented in the ArXiv version.

▶ Problem 1. Let $\alpha \geq 1$ be a constant.
Input: M, a red-blue matching.
Output: An untangle sequence starting at M of length at most α times $d(M)$.

▶ **Theorem 2.1.** *Problem 1 is NP-hard for all $\alpha \geq 1$.*

In *RPM 3-SAT*, the *graph* of a CNF formula is the bipartite graph with the variables and clauses as vertices, and where there is an edge between a variable and a clause if and only if the clause contains the variable. A CNF formula is *monotone* if each clause contains either only positive or only negative variables. An *RPM 3-CNF* formula is a monotone formula whose graph can be drawn with no intersection, and with the three following conventions (Figure 2). (i) The variables and the clauses are represented by axis-parallel rectangles. (ii) The variable rectangles lie on the x-axis. (iii) The positive clause rectangles are above the x-axis, the negative ones, below. We call such a drawing the *planar embedding* of Φ.

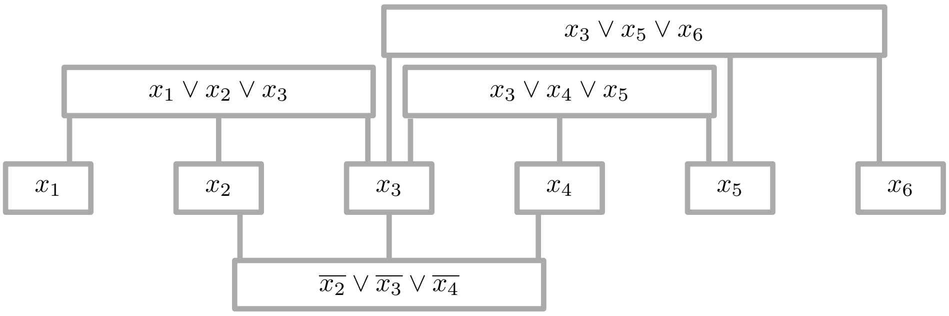

Figure 2 A planar embedding of an RPM 3-CNF formula.

The idea of the reduction is that, given an RPM 3-CNF formula Φ, we draw a rectilinear red-blue matching M_Φ of polynomial size such that all the untangle sequences starting at M_Φ are of length at most k_1 if Φ is satisfiable, and of length at least k_2 if Φ is not satisfiable.

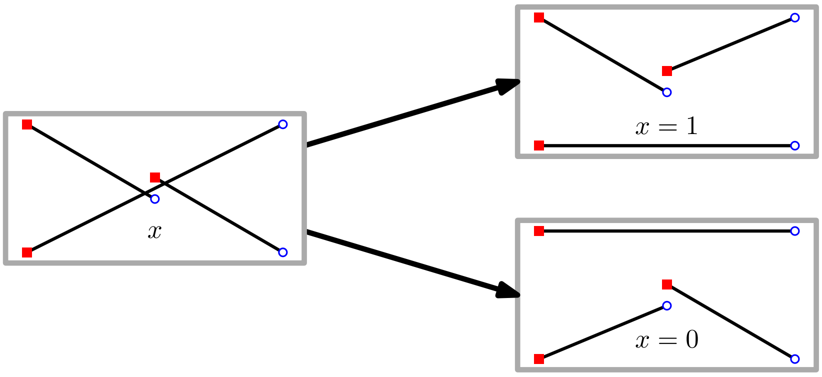

Figure 3 A variable gadget.

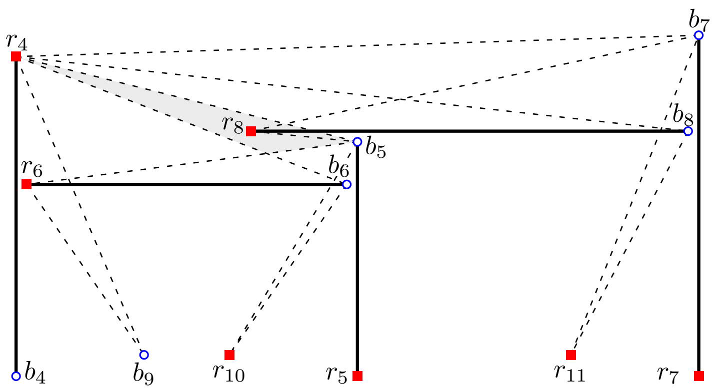

Figure 4 A clause gadget.

The aforesaid matching M_Φ is built upon the planar embedding of Φ. The variable rectangles are replaced by variable gadgets (Figure 3). The clause rectangles together with the corresponding edges are replaced with clause gadgets (Figure 4). A clause gadget consists of two OR gadgets, working like OR gates, and is connected to a padding gadget (Figure 5). If a clause is satisfied, then any untangle sequence of the two OR gadgets will end without creating any crossing in the padding gadget. If a clause is not satisfied, then any untangle sequence of the two OR gadgets will end creating a crossing in the padding gadget, which will trigger an arbitrary long series of flips, thus ensuring an arbitrary gap $k_2 - k_1$.

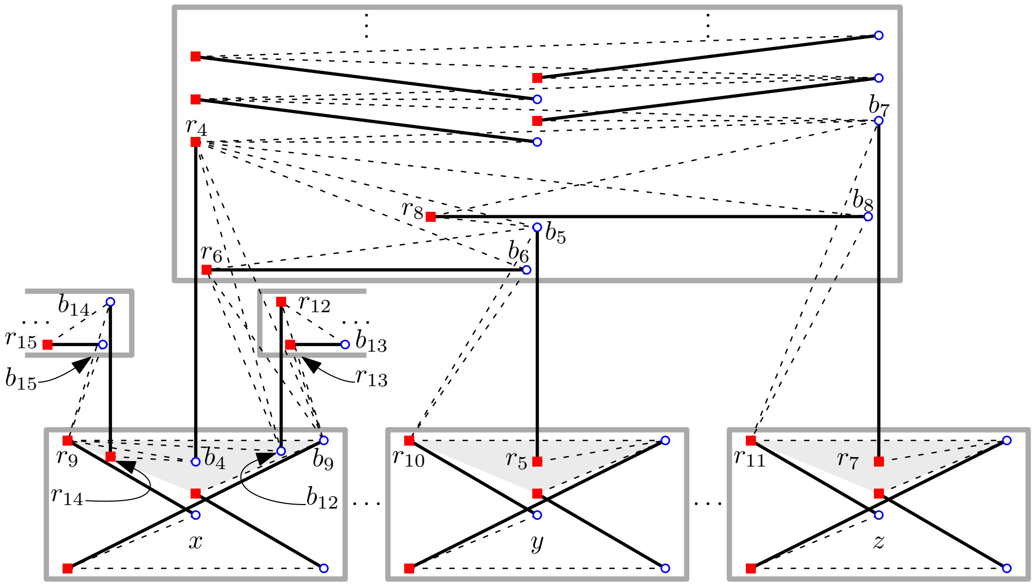

Figure 5 A clause gadget with padding connected to its variable gadgets, with branching on x.

3 Upper Bound on $\mathbf{d}(n)$

In this section, we give some insight into the proof of the following upper bound.

▶ **Theorem 3.1.** *In the red-on-a-line case,* $\mathbf{d}(n) \leq \binom{n}{2}$.

The proof consists of the analysis of the number of flips performed by the following recursive algorithm. We assume general position (no two blue points at same height). Let the *top* segment of a red-on-a-line matching be the segment with the topmost blue endpoint.

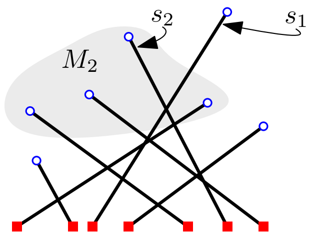

Figure 6 A red-on-a-line matching with s_1 as the top segment. The top segment of M_2 is s_2.

Algorithm 1:

Input : M, a red-on-a-line matching.
Output : An untangle sequence starting at M.

0 If $R = B = \emptyset$, then stop.

1 Let M_2 be the set of segments crossing s_1, the top segment of M (Figure 6). If M_2 is not empty, flip s_1 and s_2, the top segment of M_2, and repeat Step 0.

2 Recursively call the algorithm on the sub-matchings on both sides of the updated top segment of M.

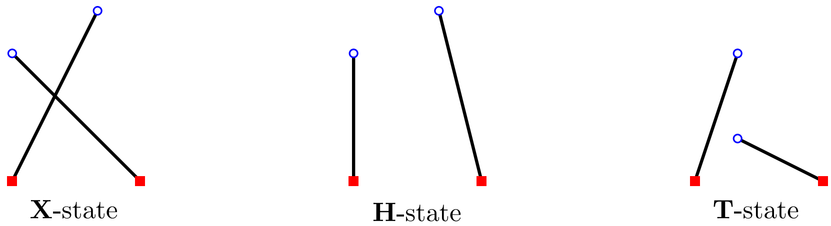

Figure 7 The three different states of pairs of segments.

The idea behind Algorithm 1 stems from the following observations. We define three states for a pair of segments: state **X**, when the segments are crossing, state **H**, when the segments are not crossing and their endpoints are in convex position, and state **T**, when the endpoints are not in convex position (Figure 7). In the convex case, a flip increases the number of **H**-pairs of at least 1 unit, providing the $\binom{n}{2}$ upper bound on $\mathbf{D}(n)$. However, the number of **H**-pairs may decrease in the general case. Figure 8 shows two such situations where there is one **H**-pair involving the segment s before the flip, and none after the flip. Algorithm 1 avoids these situations by choosing to flip top segments. The full proof, presented in the ArXiv version, involves state tracking, a novel approach to analyse flip sequences.

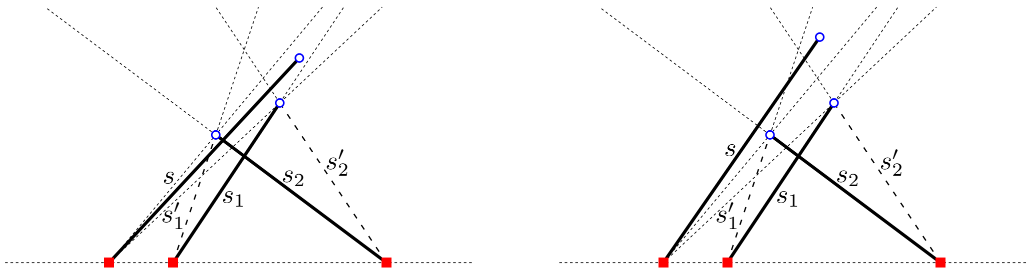

	s	s_2
s_1	**H**	**X**
s		**X**

	s	s_2'
s_1'	**X**	**H**
s		**T**

	s	s_2
s_1	**H**	**X**
s		**T**

	s	s_2'
s_1'	**T**	**H**
s		**T**

Figure 8 Two cases where the number of **H**-pairs decreases. The flipped pair is s_1, s_2.

4 Upper Bound on $\mathrm{D}(n)$

In this section, we sketch the proof of the following upper bound.

▶ **Theorem 4.1.** *In the red-on-a-line case,* $\mathbf{D}(n) \leq \binom{n}{2}\frac{n+4}{6}$.

Let $r_1, \ldots, r_n$ be the red points, ordered from left to right. Theorem 4.1 is a corollary of the following bound on the number of flips involving r_k.

▶ **Lemma 4.2.** *In the red-on-a-line case, the number of flips involving the red point r_k is at most $(k-1)(n-k)+n-1$.*

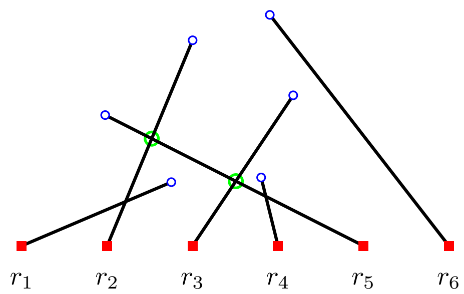

Figure 9 The two crossing pairs that may undergo a 3-flip ($k = 3$) immediately are circled.

The upper bound of Theorem 4.1 is obtained by computing the sum $\sum_{k=1}^{n}(k-1)(n-k)+n-1$ of the number of flips involving each red point, and then dividing this sum by 2, since each flip is counted twice (once for each red point).

The proof of Lemma 4.2 comes from a stronger lemma bounding the number of k-flips by $(k-1)(n-k)+n-1$, where a k-flip is a flip of a pair of segments $r_i b, r_j b'$, with $i \leq k \leq j$ (see Figure 9, where $k = 3$). The proof is fully presented in the ArXiv version.

5 Lower Bounds

In this section, we sketch the proof of the following lower bounds.

▶ **Theorem 5.1.** *In the red-on-a-line case, for even n,* $\mathbf{D}(n) \geq \frac{3}{2}\binom{n}{2} - \frac{n}{4}$.

▶ **Theorem 5.2.** *In the convex case, for n multiple of 20,* $\mathbf{d}(n) \geq 1.4 \cdot n$.

5.1 Lower Bound on $D(n)$

In order to define the starting configurations of lower bounding untangle sequences, we first provide some ad hoc definitions. We call a red-on-a-line convex matching an *n-star* when the maximum crossing number is attained, i.e. all the $\binom{n}{2}$ pairs of segments are crossing. For convenience, we say that an n-star *looks* at a point p if its blue points are all on a common line, and if p is the intersection of this line with the line on which the red points lie. We also say that two red-blue point sets R, B and R', B' are *fully crossing* if all the pairs of segments of the form $rb, r'b'$ are crossing, where $(r, b, r', b') \in R \times B \times R' \times B'$. Two matchings are fully crossing if their underlying red-blue point sets are fully crossing.

An *m-butterfly* is a red-on-a-line matching consisting of two fully crossing m-stars both looking at the same point p, where p is a median of the $2m$ red points (Figure 10). The existence of an untangle sequence of length $\frac{3}{2}\binom{2m}{2} - \frac{m}{2}$, starting at an m-butterfly is presented in the ArXiv version.

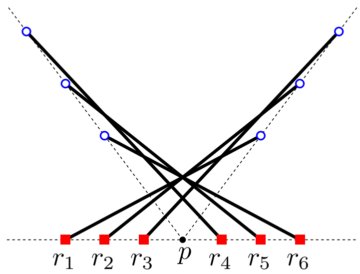

Figure 10 The 3-butterfly used to lower bound $\mathbf{D}(6)$.

5.2 Lower Bound on $\mathrm{d}(n)$

An improved lower bound on $\mathbf{d}(n)$ in the convex case comes from running a breadth-first search on the 20-segment configuration in Figure 11 and finding a minimum untangle sequence length of 24. Arranging multiple copies of this configuration, we get $\mathbf{d}(n) \geq 1.4 \cdot n$ for n multiple of 20. The source code is available on github.com/gfonsecabr/untangling.

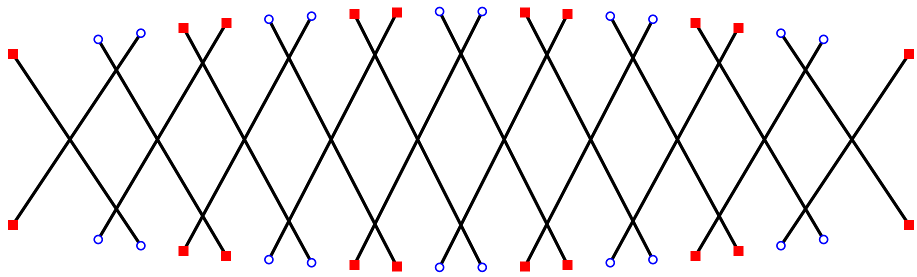

Figure 11 The convex configuration used to show that $\mathbf{d}(20) \geq 28$.

6 Concluding Remarks

Untangle sequences of TSP tours have been investigated since the 80s, when a cubic upper bound on $\mathbf{D}(n)$ has been discovered [23]. This bound also holds for matchings and has not been improved ever since. Except for the convex case, there are big gaps between the lower and upper bounds, as can be seen in Table 1. Experiments on tours and matchings have shown that, in all cases tested, the cubic upper bound is not tight and the lower bounds seem to be asymptotically tight.

Untangle sequences have many unexpected properties which make the problem harder than it seems at first sight. The following questions remain open.

1. If we add a new segment to a crossing-free matching, what is the maximum length of an untangle sequence? Notice that an $o(n^2)$ bound would lead to an $o(n^3)$ bound for $\mathbf{d}(n)$.
2. Is it always possible to find an untangle sequence that does not flip the same pair of segments twice? Using a balancing argument, we can show that the number of *distinct* flips in any untangle sequence is $O(n^{8/3})$.
3. What is the maximum number of flips involving a given point? The classic potential [23] provides a quadratic bound which leads again to $\mathbf{D}(n) = O(n^3)$.
4. Is there a potential that provides better bounds?

We proved the NP-hardness of computing the shortest untangle sequence for a red-blue matching. What is the complexity of computing the shortest untangle sequence for a TSP tour, for a red-on-a-line matching, or even for a convex instance? What about the longest untangle sequences?

References

1 Oswin Aichholzer, Wolfgang Mulzer, and Alexander Pilz. Flip distance between triangulations of a simple polygon is NP-complete. *Discrete & Computational Geometry*, 54(2):368–389, 2015.

2 Sergey Bereg and Hiro Ito. Transforming graphs with the same degree sequence. In *Computational Geometry and Graph Theory*, pages 25–32, 2008.

3 Sergey Bereg and Hiro Ito. Transforming graphs with the same graphic sequence. *Journal of Information Processing*, 25:627–633, 2017.

4 Ahmad Biniaz, Anil Maheshwari, and Michiel Smid. Flip distance to some plane configurations. *Computational Geometry*, 81:12–21, 2019.

5 Marthe Bonamy, Nicolas Bousquet, Marc Heinrich, Takehiro Ito, Yusuke Kobayashi, Arnaud Mary, Moritz Mühlenthaler, and Kunihiro Wasa. The perfect matching reconfiguration problem. In *44th International Symposium on Mathematical Foundations of Computer Science*, volume 138 of *LIPIcs*, pages 80:1–80:14, 2019.

6 Édouard Bonnet and Tillmann Miltzow. Flip distance to a non-crossing perfect matching. *Computing Research Repository*, abs/1601.05989, 2016.

7 Prosenjit Bose and Ferran Hurtado. Flips in planar graphs. *Computational Geometry*, 42(1):60–80, 2009.

8 Nicolas Bousquet and Alice Joffard. Approximating shortest connected graph transformation for trees. In *Theory and Practice of Computer Science*, pages 76–87, 2020.

9 Mark De Berg and Amirali Khosravi. Optimal binary space partitions for segments in the plane. *International Journal of Computational Geometry & Applications*, 22(03):187–205, 2012.

10 Matthias Englert, Heiko Röglin, and Berthold Vöcking. Worst case and probabilistic analysis of the 2-Opt algorithm for the TSP. *Algorithmica*, 68(1):190–264, 2014.

11 Péter L Erdős, Zoltán Király, and István Miklós. On the swap-distances of different realizations of a graphical degree sequence. *Combinatorics, Probability and Computing*, 22(3):366–383, 2013.

12 Seifollah Louis Hakimi. On realizability of a set of integers as degrees of the vertices of a linear graph. i. *Journal of the Society for Industrial and Applied Mathematics*, 10(3):496–506, 1962.

13 Seifollah Louis Hakimi. On realizability of a set of integers as degrees of the vertices of a linear graph ii. uniqueness. *Journal of the Society for Industrial and Applied Mathematics*, 11(1):135–147, 1963.

14 Ferran Hurtado, Marc Noy, and Jorge Urrutia. Flipping edges in triangulations. *Discrete & Computational Geometry*, 22(3):333–346, 1999.

15 Takehiro Ito, Erik D. Demaine, Nicholas J.A. Harvey, Christos H. Papadimitriou, Martha Sideri, Ryuhei Uehara, and Yushi Uno. On the complexity of reconfiguration problems. *Theoretical Computer Science*, 412(12):1054–1065, 2011.

16 Alice Joffard. *Graph domination and reconfiguration problems.* PhD thesis, Université Claude Bernard Lyon 1, 2020.

17 Charles L Lawson. Transforming triangulations. *Discrete Mathematics*, 3(4):365–372, 1972.

18 Anna Lubiw and Vinayak Pathak. Flip distance between two triangulations of a point set is NP-complete. *Computational Geometry*, 49:17–23, 2015.

19 Naomi Nishimura. Introduction to reconfiguration. *Algorithms*, 11(4), 2018.

20 Yoshiaki Oda and Mamoru Watanabe. The number of flips required to obtain non-crossing convex cycles. In *Kyoto International Conference on Computational Geometry and Graph Theory*, pages 155–165, 2007.

21 Alexander Pilz. Flip distance between triangulations of a planar point set is apx-hard. *Computational Geometry*, 47(5):589–604, 2014.

22 Jan van den Heuvel. The complexity of change. *Surveys in Combinatorics*, 409:127–160, 2013.

23 Jan van Leeuwen. Untangling a traveling salesman tour in the plane. In *7th Workshop on Graph-Theoretic Concepts in Computer Science*, 1981.

24 Todd G Will. Switching distance between graphs with the same degrees. *SIAM Journal on Discrete Mathematics*, 12(3):298–306, 1999.

25 Ro-Yu Wu, Jou-Ming Chang, and Jia-Huei Lin. On the maximum switching number to obtain non-crossing convex cycles. In *26th Workshop on Combinatorial Mathematics and Computation Theory*, pages 266–273, 2009.