

Untangling Segments in the Plane*

Guilherme D. da Fonseca, Yan Gerard, Bastien Rivier

Abstract

A set of n segments in the plane may form a Euclidean TSP tour, a tree, or a matching, among others. Optimal TSP tours as well as minimum spanning trees and perfect matchings have no crossing segments, but several heuristics and approximation algorithms may produce solutions with crossings. If two segments cross, then we can reduce the total length with the following flip operation. We remove a pair of crossing segments, and insert a pair of non-crossing segments, while keeping the same vertex degrees. In this paper, we consider the number of flips performed under different assumptions, using a new unifying framework that applies to tours, trees, matchings, and other types of (multi)graphs. Within this framework, we prove several new bounds that are sensitive to whether some endpoints are in convex position or not.

Keywords: Reconfiguration, Planar Geometry, Matching, Euclidean TSP

1 Introduction

In the Euclidean Traveling Salesman Problem (TSP), we are given a set P of n points in the plane and the goal is to produce a tour of minimum Euclidean length. The TSP problem, both in the Euclidean and in the more general graph versions, is one of the most studied NP-hard optimization problems, with several approximation algorithms, as well as powerful heuristics (see for example [6, 22, 27]). Multiple PTAS are known for the Euclidean version [7, 35, 39], in contrast to the general graph version that unlikely admits a PTAS [17]. It is well known that the optimal solution for the Euclidean TSP is a simple polygon, i.e., a polygon with no crossing segments. In fact, a crossing-free solution is necessary for some applications [15]. However, most approximation algorithms (including Christofides [18] and the PTAS [7, 35, 39]), as well as a variety of simple heuristics (nearest neighbor, greedy, and insertion, among others) may produce solutions with pairs of crossing segments. In practice, these algorithms may be supplemented with a local search phase, in which crossings are removed by iterative modification of the solution. The objective of this paper is to analyze this local search phase, for TSP as well as for several other problems involving planar segments, using a unifying approach. Before considering other problems, we provide some definitions and present the previous results on untangling a TSP tour.

Given a Euclidean TSP tour, a *flip* is an operation that removes a pair of crossing segments and inserts a new pair of segments preserving a tour (Figure 1(a)). In order to find a tour without crossing segments starting from an arbitrary tour, it suffices to find a crossing, perform a flip, and repeat until there are no crossings, in a process called *untangle*. The untangle always terminates as each flip strictly shortens the tour. Since a flip may create several new crossings (Figure 1(a)), it is not obvious how to bound the number of flips performed to untangle a tour. Let $\mathbf{D}_{\text{TSP}}(n)$ denote the maximum number of flips successively performed on a TSP tour with n segments.

*Preliminary versions of these results appeared in WALCOM 2023 [19], WALCOM 2024 [20], and the PhD dissertation [40]. This work is supported by the French ANR PRC grant ADDS (ANR-19-CE48-0005).




Figure 1: Examples of flips in a (a) TSP tour, (b) tree, (c) monochromatic matching, (d) multigraph, and (e) red-blue matching.

An upper bound of $\mathbf{D}_{\text{TSP}}(n) = O(n^3)$ is proved in [43], while the best known lower bound is $\mathbf{D}_{\text{TSP}}(n) = \Omega(n^2)$. In contrast, if the endpoints P satisfy the property of being in convex position, then flips always decrease the number of crossings and tight bounds of $\mathbf{D}_{\text{Convex,TSP}}(n) = \Theta(n^2)$ are easy to prove.

When the goal is to untangle a TSP tour, we are allowed to *choose* which pair of crossing segments to remove in order to perform fewer flips, which we call *removal choice*. Let $\mathbf{D}_{\text{TSP}}^{\text{R}}(n)$ denote the number of flips needed to untangle any TSP tour with n segments, assuming the best possible removal choice. If the endpoints P are in convex position, then a tight bound of $\mathbf{D}_{\text{Convex,TSP}}^{\text{R}}(n) = \Theta(n)$ has been shown [37, 46]. However, for points in general position, the best bound known is again $\mathbf{D}_{\text{TSP}}^{\text{R}}(n) \leq \mathbf{D}_{\text{TSP}}(n) = O(n^3)$, which does not use removal choice, and the only asymptotic lower bound known is the trivial $\mathbf{D}_{\text{TSP}}^{\text{R}}(n) = \Omega(n)$.

The same flip operation may be applied in other settings. More precisely, a *flip* consists of removing a pair of crossing segments s_1, s_2 and inserting a pair of segments s'_1, s'_2 in a way that s_1, s'_1, s_2, s'_2 forms a cycle and a certain *graph property* is preserved. In the case of *TSP* tours, the property is being a Hamiltonian cycle. Other properties have also been studied, such as spanning *trees* (Figure 1(b)), perfect *matchings* (Figure 1(c)), and *multigraphs* (Figure 1(d)). Notice that flips preserve the degrees of all vertices and that multiple copies of the same edge may appear when we perform a flip on certain graphs. Biniáz et al. [10] showed that if the graph property is being a tree, the points are in convex position, and we perform the correct removal choice, then $\mathbf{D}_{\text{Convex,Tree}}^{\text{R}}(n) = O(n \log n)$ flips suffice to untangle the tree.

Notice that, when the graph property is being a tour or a tree, choosing which pair of crossing edges we remove determines which pair of crossing edges we insert. However, this is not the case for matchings and multigraphs, where we are also allowed to choose which pair of segments to insert among two possibilities, which we call *insertion choice* (Figure 1(c) and (d)).

Bonnet et al. [12] showed that, using insertion choice, it is possible to untangle a matching using $\mathbf{D}_{\text{Matching}}^{\text{I}}(n) = O(n^2)$ flips. Let σ be the *spread* of P , that is, the ratio between the maximum and minimum distances among points in P . Notice that $\sigma = \Omega(\sqrt{n})$ but may be arbitrarily large. Biniáz et al. [10] showed an alternative bound of $\mathbf{D}_{\text{II,Matching}}^{\text{I}}(n, \sigma) = O(n\sigma)$ flips when the endpoints have spread σ . Biniáz et al. [10] also showed that if both removal and insertion choices are used and P is in convex position, then $\mathbf{D}_{\text{Convex,Matching}}^{\text{RI}}(n) = O(n)$ flips suffice to untangle a matching.

Another version of the problem that has been studied is the *red-blue* version (Figure 1(e)). In this version, the endpoints P are partitioned into n *red* points and n *blue* points, while the

segments must connect points of different colors, forming a red-blue perfect matching. Notice that insertion choice is not possible in this version. If the endpoints P satisfy the property Π of all red points being colinear, then using removal choice $\mathbf{D}_{\Pi, \text{RedBlue}}^{\text{R}}(n) = O(n^2)$ flips suffice [10, 21] and the only lower bound known is $\mathbf{D}_{\Pi, \text{RedBlue}}^{\text{R}}(n) = \Omega(n)$. If the endpoints P are in convex position, then a tight bound of $\mathbf{D}_{\text{Convex}, \text{RedBlue}}^{\text{R}}(n) = \Theta(n)$ is known [10].

1.1 New Results

Using removal or insertion choices to obtain shorter flip sequences has not been studied in a systematic way before and opens several new questions, while unifying the solution to multiple problems. In contrast, previous bounds are usually stated for a single graph property. For example, the seminal TSP bound of $\mathbf{D}_{\text{TSP}}(n) = O(n^3)$ flips [43] has been rediscovered (with nearly identical proofs) in the context of matchings as $\mathbf{D}_{\text{Matching}}(n) = O(n^3)$ after 35 years [12]. As another example, the tree bound of $\mathbf{D}_{\text{Convex}, \text{Tree}}^{\text{R}}(n) = O(n \log n)$ from [10] uses very specific properties of trees, while our new bound of $\mathbf{D}_{\text{Convex}}^{\text{R}}(n) = O(n \log n)$ uses a completely different approach that holds for trees, TSP tours, matchings, and arbitrary multigraphs.

Using the framework of choices, we are able to state the results in a more general setting. Upper bounds for multigraphs that do not use insertion choice apply to all aforementioned problems. In contrast, bounds that use insertion choice are unlikely to generalize to red-blue matchings, TSP tours, or trees, where insertion choice is not available. Still, they generally hold for both monochromatic matchings as well as multigraphs.

The goal of the paper is to obtain improved bounds in this framework of removal and insertion choices, providing unified proofs to multiple problems. In order to bridge the gaps between points in convex and in general position, we consider configurations where most or all points are in convex position. Let $P = C \cup T$ where C is in convex position and the points of T are placed anywhere. Let S be a multiset of n segments with endpoints P and t be the sum of the degrees of the points in T (and define $t = O(1)$ if $T = \emptyset$). We prove the following results to obtain a crossing-free solution from a set of segments S . Some bounds are summarized in Table 1.

Using no choice (Section 3): We show that a flip sequence has at most $O(tn^2)$ flips. This bound continuously interpolates between the best bounds known for the convex and general cases and holds for TSP tours, matchings, red-blue matching, trees, and multigraphs. In fact, we show that when no choice is used, all lower and upper bounds for matchings can be converted into bounds for the remaining problems (within a constant factor). We also show that any flip sequence has $O(n^{8/3})$ *distinct* flips.

Using only removal choice (Section 4): If $|T| \leq 2$, then $O(t^2n + n \log n)$ flips suffice. These bounds hold for TSP tours, matchings, red-blue matching, trees, and multigraphs. The $O(n \log n)$ term is removed for the special cases of TSP tours and red-blue matchings.

Using only insertion choice (Section 5): In the convex case ($T = \emptyset$), we show that $O(n \log n)$ flips suffice. If T is separated from C by two parallel lines, then $O(tn \log n)$ flips suffice.

Using both removal and insertion choices (Section 6): If T is separated from C by two parallel lines, then $O(tn)$ flips suffice. If T is anywhere outside the convex hull of C and S is a matching, then $O(t^3n)$ flips suffice.

These bounds improve several previous upper bounds when most points are in convex position with a small number of additional points in general position. Essentially, we obtain linear or near-linear upper bounds when we have removal or insertion choice and quadratic upper bounds

Graph:	TSP, Red-Blue		Monochromatic Matching			
Choices:	\emptyset	R	\emptyset	R	I	RI
Convex	n^2	n [10, 37, 46]	n^2	$n \log n$ (4.1)	$n \log n$ (5.1)	n [10]
$ T = 1$	n^2 (3.2)	n (4.3)	n^2 (3.2)	$n \log n$ (4.3)	n^2 [12]	n (4.3)
$ T = 2$	n^2 (3.2)	n (4.9)	n^2 (3.2)	$n \log n$ (4.9)	n^2 [12]	n (4.9)
T separated	tn^2 (3.2)				$tn \log n$ (5.2)	tn (6.1)
T outside	tn^2 (3.2)				n^2 [12]	$t^3 n$ (6.4)
T anywhere	tn^2 (3.2)				n^2 [12]	

Table 1: Upper bounds to different versions of the problem with points having $O(1)$ degree. The letter R corresponds to removal choice, I to insertion choice, and \emptyset to no choice. New results are highlighted in yellow with the theorem number in parenthesis and tight bounds are bold.

when no choice is available. Conjecture has been made that the convex upper bounds hold asymptotically for points in general position [12, 21].

In a matching or TSP tour, we have $t = O(|T|)$ and $n = O(|P|)$, however, in a tree, t can be as high as $O(n)$. In a multigraph t and n can be arbitrarily larger than $|T|$ and $|P|$. The theorems describe more precise bounds as functions of all these parameters. For simplicity, the introduction only describes the bounds in terms of n and t .

1.2 Related Reconfiguration Problems

Combinatorial reconfiguration studies the step-by-step transition from one solution to another for a given combinatorial problem. Many reconfiguration problems are presented in [41]. We give a brief overview of reconfiguration among n line segments using alternative flip operations.

The *2OPT flip* is not restricted to crossing segments. It removes and inserts pairs of segments (the four segments forming a cycle) as long as the total length decreases. In contrast to flips restricted to crossing segments, the number of 2OPT flips performed may be exponential [23].

It is possible to relax the flip definition even further to all operations that replace two segments by two others such that the four segments together form a cycle [8, 9, 11, 14, 25, 45]. This definition has also been considered for multigraphs [28, 29, 31].

Another type of flip consists of removing a single segment and inserting another one. Such flips are widely studied for triangulations [4, 30, 33, 34, 36, 38]. They have also been considered for non-crossing trees [1, 44], paths, and matchings [2]. It is possible to reconfigure any two non-crossing paths using $O(n)$ flips if the points are in convex position [5, 16, 32], if there is *one* point *inside* the convex hull [3], and if there is *one* point *outside* the convex hull [13]. It is also possible to reconfigure any two non-crossing paths (using an unknown number of flips) if the endpoints are in two convex layers [32].

1.3 Overview

In Section 2, we precisely define the notation used throughout the paper and present some key lemmas from previous work. In Section 3, we present the bounds without removal or insertion choices as well as several reductions. In Section 4, we present the bounds with removal choice. In Section 5, we present the bounds with insertion choice. In Section 6, we present the bounds with both removal and insertion choices. Concluding remarks and open problems are presented in Section 7.

2 Preliminaries

This section contains a collection of definitions, and lemmas used throughout the paper.

Asymptotic bounds incorrectly fading to zero. Our bounds often have terms like $O(tn)$ and $O(n \log |C|)$ that would incorrectly become 0 if t or $\log |C|$ is 0. In order to avoid this problem, factors in the O notation should be made at least 1. For example, the aforementioned bounds should be respectively interpreted as $O((1+t)n)$ and $O(n \log(2 + |C|))$.

General position. Throughout this paper, we assume general position of the endpoints of the segments we consider. In particular, we exclude three colinear endpoints, in which case a flip and a crossing are not well defined. In some proofs, we assume that all y -coordinates are distinct, but it is straightforward to remove this assumption. We also assume that the two endpoints of a segment are always distinct.

2.1 General Definitions

In the following, we summarize important definitions used throughout this paper.

Segment types. Given two (possibly equal) sets P_1, P_2 of endpoints, we say that a segment is a P_1P_2 -segment if one of its endpoints is in P_1 and the other is in P_2 .

Crossings. We say that two segments *cross* if they intersect at a single point which is not an endpoint of either segment. A pair of crossing segments is called a *crossing*. We also say that a line and a segment *cross* if they intersect at a single point which is not an endpoint of the segment.

If h is a segment or a line (respectively a set of two parallel lines) we say that h *separates* a set of points P if P can be partitioned into two non-empty sets P_1, P_2 such that every segment p_1p_2 with $p_1 \in P_1, p_2 \in P_2$ crosses h (respectively crosses at least one of the lines in h).

2.2 Properties

We freely interpret a multiset S of n segments with endpoints P in the plane as the multigraph (P, S) where the set of vertices is the set of endpoints and where the multiset of edges is the multiset of segments. In some cases the endpoints P have a color (typically red or blue).

We consider the following two types of properties.

1. The *point properties* are the properties about the position of the endpoints. A notable point property is that the endpoints are in convex position. Throughout, we refer to this property as **Convex**.
2. The *graph properties* are the class of graphs before and after a flip. Example: the segments form a matching.

The **graph properties** Γ considered in this paper are the following:

- **TSP:** Being a Hamiltonian cycle.
- **Matching:** Being a perfect matching.
- **RedBlue:** Being a perfect matching where each segment connects a red endpoint to a blue endpoint.

- **Tree:** Being a spanning tree.
- **Multigraph:** Being a multigraph. Since this property imposes no restriction on the segments, it may not be explicitly written on the bounds.

2.3 Flips and Flip Graphs

Given a graph property Γ and a multiset of segments S satisfying Γ , a *flip* f is the operation of removing two crossing segments $s_1, s_2 \in S$ and inserting two non-crossing segments s'_1, s'_2 with the same four endpoints that preserve the property Γ .

Given a set of endpoints P , an integer n , and a graph property Γ , the *flip graph* $F(P, n, \Gamma)$ is defined as the following directed acyclic simple graph. The vertices of the flip graph are the multisets of n segments with endpoints in P that satisfy the property Γ . There is a directed arc from S_1 to S_2 if there exists a flip f such that $S_2 = f(S_1)$. Figure 2(a) and 3(a) display examples of flip graphs in the TSP and matching versions, respectively.




Figure 2: The (a) flip graph and a (b) removal strategy of a five-point set in the TSP version.

Given a flip graph F , a *strategy* is a spanning subgraph of F with the same set of sink vertices (a sink vertex corresponds to a crossing-free set of segments). A *removal strategy* is a strategy R such that each vertex S of F has an associated crossing pair of segments $s_1, s_2 \in S$ and R only contains the arcs from S that remove s_1, s_2 . Figures 2(b) and 3(b) display examples of removal strategies in the TSP and matching versions, respectively.

The following strategies with insertion choice are only possible if the graph property Γ is such that for every crossing pair of segments removed, there are two possible pairs of segment that may be added. A *strategy for both removal and insertion choices* is a strategy RI such that each vertex of RI that is not a sink in F has out-degree exactly one. An *insertion strategy* is a strategy I obtained from F by erasing the following arcs. For each pair of arcs f, f' coming from the same vertex S and with the same pair of crossing segments $s_1, s_2 \in S$ removed, exactly one of the two arcs f, f' is erased and the other is kept in I . Figure 3(c) and 3(d) display examples of a removal strategy and a strategy for both choices, respectively, in the matching version.

Let $\mathcal{R}(F), \mathcal{I}(F), \mathcal{RI}(F)$ respectively denote the collection of all strategies for the flip graph F that are removal strategies, insertion strategies, and strategies for both removal and insertion choices.

A *flip sequence* is a path in the flip graph (or a strategy, which is a spanning subgraph of the flip graph) and an *untangle sequence* is a flip sequence that ends in a sink of the flip graph. The




Figure 3: (a) The flip graph F of a six-point set in the matching version. (b) Removal, (c) insertion, and (d) both removal and insertion strategies for the flip graph F . Pairs of edges corresponding to the same crossing are grouped.

length of a flip sequence is the number of flips (the arcs of the flip graph) it contains. Given a graph F , let $L(F)$ denote the length of the longest directed path in F . Notice that $L(F) = 3$ for the flip graph F in Figure 3(a), $L(R) = 2$ for the removal strategy R in Figure 3(b), $L(I) = 2$ for the insertion strategy I in Figure 3(c), and $L(RI) = 1$ for the removal and insertion strategy RI in Figure 3(d).

We are now ready to formally define the following four parameters, which are the central object of this paper.

$$\begin{aligned}
\mathbf{D}_{\Pi,\Gamma}(n) &= \max_{P \text{ satisfying } \Pi} L(F(P, n, \Gamma)) \\
\mathbf{D}_{\Pi,\Gamma}^{\mathbf{R}}(n) &= \max_{P \text{ satisfying } \Pi} \min_{R \in \mathcal{R}(F(P, n, \Gamma))} L(R) \\
\mathbf{D}_{\Pi,\Gamma}^{\mathbf{I}}(n) &= \max_{P \text{ satisfying } \Pi} \min_{I \in \mathcal{I}(F(P, n, \Gamma))} L(I) \\
\mathbf{D}_{\Pi,\Gamma}^{\mathbf{RI}}(n) &= \max_{P \text{ satisfying } \Pi} \min_{RI \in \mathcal{RI}(F(P, n, \Gamma))} L(RI)
\end{aligned}$$

To simplify the notation, we omit Π when it is the most general point property of being a set of points in general position, and we omit Γ when it is the most general graph property of being any multigraph.

It is convenient to imagine that two clever players, let us call them *the oracle* and *the adversary*, are playing the following game using the flip graph as a board. The oracle aims at minimizing the number of flips, while the adversary aims at maximizing it. Initially, the adversary chooses the starting vertex in the flip graph. The game ends when a crossing-free vertex is reached. The four definitions corresponds to the number of flips with different choices for the oracle (specified in the exponent). The number of flips is \mathbf{D} if the adversary performs all the choices, it is $\mathbf{D}^{\mathbf{R}}$ if the oracle performs all the removal choices while the adversary performs all the insertion choices, it is $\mathbf{D}^{\mathbf{I}}$ if the oracle performs all the insertion choices while the adversary performs all the removal choices, it is $\mathbf{D}^{\mathbf{RI}}$ if the oracle performs all the removal and all the insertion choices (the adversary only gets to choose the starting vertex).

2.4 Reductions

In this section, we state some trivial inequalities between the different versions of \mathbf{D} . First, as we have fewer choices, the value of \mathbf{D} can only increase.

Lemma 2.1. *The following inequalities hold for any non-negative integer n , point property Π , and graph property Γ .*

$$\begin{aligned}
\mathbf{D}_{\Pi,\Gamma}^{\mathbf{RI}}(n) &\leq \mathbf{D}_{\Pi,\Gamma}^{\mathbf{R}}(n) \leq \mathbf{D}_{\Pi,\Gamma}(n) \\
\mathbf{D}_{\Pi,\Gamma}^{\mathbf{RI}}(n) &\leq \mathbf{D}_{\Pi,\Gamma}^{\mathbf{I}}(n) \leq \mathbf{D}_{\Pi,\Gamma}(n)
\end{aligned}$$

Also, as the point set becomes less constrained, the value of \mathbf{D} can only increase. Furthermore, if the graph property is less constrained, it is easy to show that the value of \mathbf{D} can only increase using the fact that insertion strategies are only possible when both insertion choices preserve the graph property at every flip.

Lemma 2.2. *The following inequality holds for any non-negative integer n , two point properties Π, Π' such that Π implies Π' , two graph properties Γ, Γ' such that Γ implies Γ' , and for any Choices $\in \{\emptyset, \mathbf{R}, \mathbf{I}, \mathbf{RI}\}$ that are possible for Γ .*

$$\mathbf{D}_{\Pi,\Gamma}^{\text{Choices}}(n) \leq \mathbf{D}_{\Pi',\Gamma'}^{\text{Choices}}(n)$$

Finally, we present some inequalities relating matchings and red-blue matchings.

Lemma 2.3. *The following inequalities hold for any non-negative integer n , and for any point property Π .*

$$\begin{aligned} \mathbf{D}_{\Pi, \text{Matching}}^{\text{RI}}(n) &\leq \mathbf{D}_{\Pi, \text{RedBlue}}^{\text{R}}(n) \\ \mathbf{D}_{\Pi, \text{Matching}}^{\text{I}}(n) &\leq \mathbf{D}_{\Pi, \text{RedBlue}}(n) \end{aligned}$$

2.5 Splitting Lemma

The notion of splitting first appears in [12]. The idea is that some subsets of segments may be untangled independently. We start with a definition.

Given a multiset S of segments, a *splitting partition* is a partition of S into $k > 1$ subsets S_1, \dots, S_k such that the following condition holds. Let P_i be the set of endpoints of S_i and $\binom{P_i}{2}$ be the set of segments with endpoints in P_i . The condition is that there is no pair of crossing segments $s_i \in \binom{P_i}{2}$ and $s_j \in \binom{P_j}{2}$ for $i \neq j$.

Lemma 2.4 ([12, 20]). *Consider a multiset of segments S and a splitting partition S_1, \dots, S_k . The subsets S_1, \dots, S_k may be untangled independently in any order to obtain an untangle sequence of S and every untangle sequence of S may be obtained in such a way.*

If a singleton $\{s\}$ is in a splitting partition, then we say that the segment s is *uncrossable*.

3 Untangling with No Choice

In this section, we study \mathbf{D} , the length of the longest untangle sequence in the flip graph. First, we prove reductions between several versions of \mathbf{D} . Then, we prove an upper bound parameterized by t (the sum of the degrees of the points which are not in convex position). Finally, we prove a sub-cubic upper bound on the number of *distinct* flips in an untangle sequence. This number is *not* the length of an untangle sequence where the same flip may appear several times (up to a linear number of times) and is counted with its multiplicity.

3.1 Reductions

In this section, we provide a series of inequalities relating the different versions of \mathbf{D} . In particular, we show that $\mathbf{D}_{\text{Matching}}$, $\mathbf{D}_{\text{RedBlue}}$, \mathbf{D}_{TSP} , \mathbf{D}_{Tree} , and $\mathbf{D}_{\text{Multigraph}}$ have the same asymptotic behavior.

We say that a point property Π is *replicable* if for any point set P that satisfies Π , any function $k : P \rightarrow \mathbb{N}$, and any $\epsilon > 0$, there exists a point set P' that satisfies Π replacing each point $p \in P$ by $k(p)$ points located within distance at most ϵ from p . Notice that point properties of being in general and convex position are replicable, while the point property of having at most one point inside the convex hull is not.

Theorem 3.1. *For all positive integer n and for all replicable point property Π , we have the following relations:*

$$\begin{aligned} \mathbf{D}_{\Pi, \text{Matching}}(n) &= \mathbf{D}_{\Pi, \text{Multigraph}}(n), & (1) \\ 2\mathbf{D}_{\Pi, \text{Matching}}(n) &\leq \mathbf{D}_{\Pi, \text{RedBlue}}(2n) \leq \mathbf{D}_{\Pi, \text{Matching}}(2n), & (2) \\ 2\mathbf{D}_{\Pi, \text{RedBlue}}(n) &\leq \mathbf{D}_{\Pi, \text{TSP}}(3n) \leq \mathbf{D}_{\Pi, \text{Matching}}(3n), & (3) \\ 2\mathbf{D}_{\Pi, \text{RedBlue}}(n) &\leq \mathbf{D}_{\Pi, \text{Tree}}(3n) \leq \mathbf{D}_{\Pi, \text{Matching}}(3n). & (4) \end{aligned}$$

Proof. Equality (1) can be rewritten $\mathbf{D}_{\Pi, \text{Multigraph}}(n) \leq \mathbf{D}_{\Pi, \text{Matching}}(n) \leq \mathbf{D}_{\Pi, \text{Multigraph}}(n)$. Hence, we have to prove eight inequalities. The right-side inequalities are trivial, since the left-side property is stronger than the right-side property (using the equality (1) for inequalities (3) and (4)).

The proofs of the remaining inequalities follow the same structure: given a flip sequence of the left-side version, we build a flip sequence of the right-side version, having similar length and number of points.

Proving $\mathbf{D}_{\Pi, \text{Multigraph}}(n) \leq \mathbf{D}_{\Pi, \text{Matching}}(n)$ (1). We prove the left inequality of (1). A point of degree deg larger than 1 can be replicated as deg points that are arbitrarily close to each other in order to produce a matching of $2n$ points. This replication preserves the crossing pairs of segments, possibly creating new crossings (Figure 4(a)). Thus, for any flip sequence in the **Multigraph** version, there exists a flip sequence in the **Matching** version of equal length, yielding $\mathbf{D}_{\Pi, \text{Multigraph}}(n) \leq \mathbf{D}_{\Pi, \text{Matching}}(n)$.




Figure 4: (a) A crossing-free multigraph transformed into a matching by replacing multi-degree points by clusters of points. (b) Two flips in the **RedBlue** version simulating one flip in the **Matching** version. (c) Two flips in the **TSP** version simulating one flip in the **RedBlue** version. (d) Two flips in the **Tree** version simulating one flip in the **RedBlue** version.

Proving $2\mathbf{D}_{\Pi, \text{Matching}}(n) \leq \mathbf{D}_{\Pi, \text{RedBlue}}(2n)$ (2). The left inequality of (2) is obtained by duplicating the monochromatic points of the matching S into two arbitrarily close points, one red and the other blue. Then each segment of S is also duplicated into two red-blue segments. We obtain a red-blue matching S' with $2n$ segments. A crossing in S corresponds to four crossings in S' . Flipping this crossing in S amounts to choose which of the two possible pairs of segments replaces the crossing pair. It is simulated by flipping the two crossings in S' such that the resulting pair of double segments corresponds to the resulting pair of segments of the initial flip. These two crossings always exist and it is always possible to flip them one after the other as they involve disjoint pairs of segments. Figure 4(b) shows this construction. A sequence of k flips on S provides a sequence of $2k$ flips on S' . Hence, $2\mathbf{D}_{\Pi, \text{Matching}}(n) \leq \mathbf{D}_{\Pi, \text{RedBlue}}(2n)$.

Proving $2\mathbf{D}_{\Pi, \text{RedBlue}}(n) \leq \mathbf{D}_{\Pi, \text{TSP}}(3n)$ (3). To prove the left inequality of (3), we start from a red-blue matching S with $2n$ points and n segments and build a cycle S' with $3n$ points and $3n$ segments. We then show that the flip sequence of length k on S provides a flip sequence of length $2k$ on S' . We build S' in the following way. Given a red-blue segment $rb \in S$, the red point r is duplicated in two arbitrarily close points r and r' which are adjacent to b in S' . We still need to connect the points r and r' in order to obtain a cycle S' . We define S' as the cycle $r_1, b_1, r'_1, \dots, r_i, b_i, r'_i, \dots, r_n, b_n, r'_n, \dots$ where r_i is matched to b_i in S (Figure 4(c)).

We now show that a flip sequence of S with length k provides a flip sequence of S' with length $2k$. For a flip on S removing the pair of segments $r_i b_i, r_j b_j$ and inserting the pair of segments $r_i b_j, r_j b_i$, we perform the following two successive flips on S' .

- The first flip removes $r_i b_i, r'_j b_j$ and inserts $r_i b_j, r'_j b_i$.
- The second flip removes $r'_i b_i, r_j b_j$ and inserts $r'_i b_j, r_j b_i$.

The cycle then becomes $r_1, b_1, r'_1, \dots, r_i, b_j, r'_j, \dots, r_j, b_i, r'_i, \dots, r_n, b_n, r'_n, \dots$ on which we can apply the next flips in the same way. Hence, $2\mathbf{D}_{\Pi, \text{Matching}}(n) \leq \mathbf{D}_{\Pi, \text{TSF}}(3n)$.

Proving $2\mathbf{D}_{\Pi, \text{RedBlue}}(n) \leq \mathbf{D}_{\Pi, \text{Tree}}(3n)$ (4). The proof of the left inequality of (4) follows the exact same construction as in the proof of the left inequality of (3). The only difference is that, in order for S' to form a tree and not a cycle, we omit the segment $r'_n r_1$, yielding a polygonal line which is a tree (Figure 4(d)). \square

3.2 Upper Bound for Near Convex Position

In this section, we bridge the gap between the $O(n^2)$ bound on the length of untangle sequences for a set P of points in convex position and the $O(n^3)$ bound for P in general position. We prove the following theorem in the **Matching** version; the translation to the multigraph version, follows from the reductions in Theorem 3.1 and then other graph properties follow from Lemma 2.2.

Theorem 3.2. *Consider a multiset S of n segments with endpoints P partitioned into $P = C \cup T$ where C is in convex position. Let t be the sum of the degrees of the points in T . Any untangle sequence of S has length at most*

$$\mathbf{D}(n, t) = O(tn^2).$$

Proof. The proof strategy is to combine the potential χ used in [10] with the potential Λ_L used in [42]. Given a matching S , the potential $\chi(S)$ is defined as the number of crossing pairs of segments in S . Since there are n segments in S , $\chi(S) \leq \binom{n}{2} = O(n^2)$. Unfortunately, with points in non-convex position, a flip f might *increase* (or leave unchanged) χ , i.e. $\chi(f(S)) \geq \chi(S)$.

The potential Λ_L is derived from the line potential introduced in [43] but instead of using the set of all the $O(n^2)$ lines through two points of P , we use a subset of $O(tn)$ lines in order to take into account that only t points are in non-convex position. More precisely, let the potential $\Lambda_\ell(S)$ of a line ℓ be the number of segments of S crossing ℓ . Note that $\Lambda_\ell(S) \leq n$. The potential $\Lambda_L(S)$ is then defined as follows: $\Lambda_L(S) = \sum_{\ell \in L} \Lambda_\ell(S)$.

We now define the set of lines L as the union of L_1 and L_2 , defined hereafter. Let C be the subset containing the $2n - t$ points of P which are in convex position. Let L_1 be the set of the $O(tn)$ lines through two points of P , at least one of which is not in C . Let L_2 be the set of the $O(n)$ lines through two points of C which are consecutive on the convex hull boundary of C .

Let the potential $\Phi(S) = \chi(S) + \Lambda_L(S)$. We have the following bounds: $0 \leq \Phi(S) \leq O(tn^2)$. To complete the proof of Theorem 3.2, we show that any flip decreases Φ by at least 1.

We consider an arbitrary flip f removing the pair of segments $p_1 p_3, p_2 p_4$ and inserting the pair of segments $p_1 p_4, p_2 p_3$. Let o be the point of intersection of $p_1 p_3$ and $p_2 p_4$. It is shown in [43] that f never increases the potential Λ_ℓ of a line ℓ . More precisely, we have the following three cases:

- The potential Λ_ℓ decreases by 1 if the line ℓ separates the final segments $p_1 p_4$ and $p_2 p_3$ and exactly one of the four flipped points belongs to ℓ . We call these lines *f-critical* (Figure 5(a)).
- The potential Λ_ℓ decreases by 2 if the line ℓ strictly separates the final segments $p_1 p_4$ and $p_2 p_3$. We call these lines *f-dropping* (Figure 5(b)).
- The potential Λ_ℓ remains stable in the remaining cases.




Figure 5: (a) An f -critical line ℓ for a flip f removing p_1p_3, p_2p_4 and inserting p_1p_4, p_2p_3 . This situation corresponds to case (2a) with $\ell \in L_1$. (b) An f -dropping line ℓ . This situation corresponds to case (2b) with $\ell \in L_2$.

Notice that, if a point q lies in the triangle p_1op_4 , then the two lines qp_1 and qp_4 are f -critical (Figure 5(a)).

To prove that Φ decreases, we have the following two cases.

Case 1. If χ decreases, as the other term Λ_L does not increase, then their sum Φ decreases as desired.

Case 2. If not, then χ increases by an integer k with $0 \leq k \leq n - 1$, and we know that there are $k + 1$ new crossings after the flip f . Each new crossing involves a distinct segment with one endpoint, say q_i ($0 \leq i \leq k$), inside the non-simple polygon p_1, p_4, p_2, p_3 (Figure 5). Next, we show that each point $q \in \{q_0, \dots, q_k\}$ maps to a distinct line in L which is either f -dropping or f -critical, thus proving that the potential Λ_L decreases by at least $k + 1$.

We assume without loss of generality that q lies in the triangle p_1op_4 . We consider the two following cases.

Case 2a. If at least one among the points q, p_1, p_4 is not in C , then either qp_1 or qp_4 is an f -critical line $\ell \in L_1$ (Figure 5(a)).

Case 2b. If not, then q, p_1, p_4 are all in C , and the two lines through q in L_2 are both either f -dropping (the line ℓ in Figure 5(b)) or f -critical (the line qp_4 in Figure 5(b)). Consequently, there are more lines $\ell \in L_2$ that are either f -dropping or f -critical than there are such points $q \in C$ in the triangle p_1op_4 , and the theorem follows. \square

3.3 Upper Bound without Multiplicity

In this section, we prove the following theorem. We recall that two flips are *distinct* if the set of the two removed and the two inserted segments of one flip is distinct from the set of the other flip.

Theorem 3.3. *Consider a multiset S of n segments with endpoints P . Any untangle sequence of S has $O(n^{8/3})$ distinct flips.*

The proof of Theorem 3.3 is based on a balancing argument from [24] and is decomposed into two lemmas that consider a flip f and two sets of segments S and $S' = f(S)$. Similarly to [43], let L be the set of lines defined by all pairs of points in $\binom{P}{2}$. For a line $\ell \in L$, let $\Lambda_\ell(S)$ be the number of segments of S crossed by ℓ and $\Lambda_L(S) = \sum_{\ell \in L} \Lambda_\ell(S)$. Notice that $\Lambda_L(S) - \Lambda_L(S')$ depends only on the flip f . The following lemma follows immediately from the fact that $\Lambda_L(S)$ takes integer values between 0 and $O(n^3)$.

Lemma 3.4. *For any integer k , the number of flips f in a flip sequence with $\Lambda_L(S) - \Lambda_L(S') \geq k$ is $O(n^3/k)$.*

Lemma 3.4 bounds the number of flips (distinct or not) that produce a large potential drop in a flip sequence. Next, we bound the number of distinct flips that produce a small potential

drop. The bound considers all possible flips on a fixed set of points and does not depend on a particular flip sequence.

Lemma 3.5. *For any integer k , the number of distinct flips f with $\Lambda_L(S) - \Lambda_L(S') < k$ is $O(n^2k^2)$.*

Proof. Let F be the set of flips with $\Lambda_L(S) - \Lambda_L(S') < k$ where $S' = f(S)$. We need to show that $|F| = O(n^2k^2)$. Consider a flip $f \in F$ removing the pair of segments p_1p_3, p_2p_4 and inserting the pair of segments p_1p_4, p_2p_3 . Next, we show that there are at most $4k^2$ such flips with a fixed final segment p_1p_4 . Since there are $O(n^2)$ possible values for p_1p_4 , the lemma follows. We show only that there are at most $2k$ possible values for p_3 . The proof that there are at most $2k$ possible values for p_2 is analogous.




Figure 6: Illustration for the proof of Lemma 3.5.

We sweep the points in $P \setminus \{p_4\}$ by angle from the ray p_1p_4 . As shown in Figure 6, let q_1, \dots, q_k be the first k points produced by this sweep in one direction, q_{-1}, \dots, q_{-k} in the other direction and $Q = \{q_{-k}, \dots, q_{-1}, q_1, \dots, q_k\}$. To conclude the proof, we show that p_3 must be in Q . Suppose $p_3 \notin Q$ for the sake of a contradiction and assume without loss of generality that p_3 is on the side of q_i with positive i . Then, consider the lines $L' = \{p_1q_1, \dots, p_1q_k\}$. Notice that $L' \subseteq L$, $|L'| = k$, and for each $\ell \in L'$ we have $\Lambda_\ell(S) > \Lambda_\ell(S')$, which contradicts the hypothesis that $\Lambda_L(S) - \Lambda_L(S') < k$. \square

Theorem 3.3 is a consequence of Lemmas 3.4 and 3.5 with $k = n^{1/3}$.

4 Untangling with Removal Choice

In this section, we devise strategies for removal choice to untangle multisets of segments with endpoints $P = C \cup T$ (where C is in convex position), therefore providing upper bounds for several versions of \mathbf{D}^R . Recall that such removal strategies choose which pair of crossing segments is removed, *but not* which pair of segments with the same endpoints is subsequently inserted. We start with a point set in convex position, followed by 1 point inside or outside the convex, then 1 point inside and 1 outside the convex, 2 points inside the convex, and 2 points outside the convex. As only removal choice is used, all results also apply to all versions.

4.1 Upper Bounds for Convex Position

Let $P = C = \{p_1, \dots, p_{|C|}\}$ be a set of points in convex position sorted in counterclockwise order along the convex hull boundary and consider a set of segments S with endpoints P . Given a segment p_ap_b and assuming without loss of generality that $a < b$, we define the *crossing depth*

$\delta_\times(p_ap_b)$ as the number of points in p_{a+1}, \dots, p_{b-1} that are an endpoint of a segment in S that crosses any other segment in S (not necessarily p_ap_b). We use the crossing depth to prove an $O(n \log n)$ bound in the convex **Multigraph** version.




Figure 7: Proof of Theorem 4.1. (a) The segments of a convex multigraph are labeled with the crossing depth. (b,c) Two possible pairs of inserted segments, with one segment of the pair having crossing depth $\lfloor \frac{3}{2} \rfloor = 1$.

Theorem 4.1. *Consider a multiset S of n segments with endpoints $P = C$ in satisfying the property **Convex**. There exists a removal strategy R such that any untangle sequence of R has length at most*

$$\mathbf{D}_{\text{Convex}}^R(n) = O(n \log |C|) = O(n \log n).$$

Proof. We repeat the following procedure until there are no more crossings. Let $p_ap_b \in S$ be a segment with crossings (hence, crossing depth at least one) and $a < b$ minimizing $\delta_\times(p_ap_b)$ (Figure 7(a)). Let $q_1, \dots, q_{\delta_\times(p_ap_b)}$ be the points defining $\delta_\times(p_ap_b)$ in order and let $i = \lceil \delta_\times(p_ap_b)/2 \rceil$. Since p_ap_b has minimum crossing depth, the point q_i is the endpoint of segment $q_i p_c$ that crosses p_ap_b . When flipping $q_i p_c$ and p_ap_b , we obtain a segment s (either $s = q_i p_a$ or $s = q_i p_b$) with $\delta_\times(s)$ at most half of the original value of $\delta_\times(p_ap_b)$ (Figure 7(b,c)). Hence, this operation always divides the value of the smallest positive crossing depth by at least two. As the crossing depth is an integer smaller than $|C|$, after performing this operation $O(\log |C|)$ times, it produces a segment of crossing depth 0. As the segments of crossing depth 0 can no longer participate in a flip, the claimed bound follows. \square

4.2 Upper Bound for One Point Inside or Outside a Convex

In this section, we prove an upper bound on the number of flips to untangle a multiset of segments with all but one endpoint in convex position. We first state a lemma used to prove Theorem 4.3.

Lemma 4.2. *Consider a set C of points in convex position, and a multiset S of n crossing-free segments with endpoints in C . Consider the multiset $S \cup \{s\}$ where s is an extra segment with one endpoint in C and one endpoint q anywhere in the plane. There exists a removal strategy R such that any untangle sequence in R starting starting from $S \cup \{s\}$ has length at most $n - 1 = O(n)$.*

Proof. Iteratively flip the segment $s = qp_1$ with the segment $p_2 p_3 \in S$ crossing qp_1 the farthest from q (Figure 8). This flip inserts a CC -segment, say $p_1 p_2$, which is impossible to flip again, because the line $p_1 p_2$ is crossing free. The flip does not create any crossing between CC -segments of the multiset S . The last flip inserts two crossing-free segments instead of one, hence the -1 in the lemma statement. \square

We are now ready to state and prove the theorem.



Figure 8: Proof of Lemma 4.2 with q (a) inside, and (b) outside the convex hull of C .

Theorem 4.3. Consider a multiset S of n segments such that the endpoints P satisfy the property Π of being partitioned into $P = C \cup T$ where C is in convex position, and $T = \{q\}$. Let t be the sum of the degrees of the points in T . Let $\mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(n)$ be the number of flips to untangle any multiset of at most n segments with endpoints in convex position, a graph property Γ , and removal choice. There exists a removal strategy R such that any untangle sequence in R has length at most

$$\mathbf{D}_{\Pi, \Gamma}^{\mathbf{R}}(n) \leq t \cdot (n - 1) + \mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(n) = O(tn + \mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(n)).$$

In particular, we have the following upper bounds for different graph properties:

$$\mathbf{D}_{\Pi}^{\mathbf{R}}(n, t) = O(n \log |C| + tn) = O(n \log n + tn) \text{ and}$$

$$\mathbf{D}_{\Pi, \text{TSP}}^{\mathbf{R}}(n, t), \mathbf{D}_{\Pi, \text{RedBlue}}^{\mathbf{R}}(n, t) = O(tn).$$

Furthermore, if S has no pair of CC -segments that cross each other, then any untangle sequence in R starting from S has length at most $t \cdot (n - 1) = O(tn)$.




Figure 9: Proof of Theorem 4.3. (a) Case 1. (b) Case 2. (c) Case 3. The existence of the the gray segment is assumed for a contradiction.

Proof. We start by untangling the segment with both endpoints in C using $\mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(n)$ flips. For each segment s with endpoint q with crossing, we apply Lemma 4.2 to s and the CC -segments in S crossing s . Once a segment s incident to q is crossing free, it is impossible to flip it again as we fall in one of the following cases (Figure 9). Let ℓ be the line containing s .

Case 1: If ℓ is crossing free, then ℓ splits (see Lemma 2.4) the multigraph in three partitions: the segments on one side of ℓ , the segments on the other side of ℓ , and the segment s itself.

Case 2: If ℓ is not crossing free and q is outside the convex hull of C , then s is uncrossable (see the splitting lemma, i.e., Lemma 2.4).

Case 3: If q is inside the convex hull of C , then introducing a crossing on s would require that q lies in the interior of the convex quadrilateral (shaded in yellow in Figure 9(c)) whose diagonals are the two segments removed by a flip. The procedure excludes this possibility by ensuring that there are no crossing pair of CC -segments in S , and, therefore, that one of the removed segment already has q as an endpoint.

Therefore, we need at most $n - 1$ flips for each of the t segments incident to q . \square

4.3 Upper Bound for One Point Inside and One Point Outside a Convex

Given an endpoint p , let $\deg(p)$ denote the degree of p , that is, the number of segments incident to p . The following lemma is used to prove Theorem 4.5.

Lemma 4.4. *Consider a multiset S of n segments with endpoints P partitioned into $P = C \cup T$ where C is in convex position, and $T = \{q, q'\}$ such that q is outside the convex hull of C and q' is inside the convex hull of C . Consider that q is the endpoint of a single segment s in S and that all the crossings of S are on s . We define the parameter t as the sum of the degrees of the points in T , i.e., $t = \deg(q') + 1$.*

There exists a removal strategy R such that any flip sequence of R starting at S and ending in a multiset of segments where all crossings (if any) are on the segment qq' (if $qq' \notin S$ then there are no crossings) has length $O(tn)$.

Proof. We proceed as follows, while s has crossings. For induction purpose, let $f(n)$ be the length of the flip sequence in the lemma statement for n segments. Let s' be the segment that crosses s at the point farthest from q . We flip s and s' , arriving at one of the three cases below (Figure 10).




Figure 10: The three cases in the proof of Lemma 4.4.

Case 1 (CT \times CC). In this case, the segment s' is a CC -segment. Notice that the line ℓ containing s' becomes crossing free after the flip. There are segments on both sides of ℓ . If ℓ separates q, q' , then we untangle both sides independently (see Lemma 2.4) using $O(n)$ and $O(tn)$ flips (Theorem 4.3). Otherwise, the segments on one side of ℓ are already crossing free (because of the specific choice of s') and we inductively untangle the $n' \leq n - 1$ segments on the other side of ℓ using $f(n')$ flips.

Case 2 (CT \times CT \rightarrow CC, TT). If s' is a CT -segment and one of the inserted segments is the TT -segment qq' , then the procedure is over as all crossings are on qq' .

Case 3 (CT \times CT \rightarrow CT, CT). In this case two CT -segments are inserted. Let $p \in C$ be an endpoint of $s = qp$. Since the inserted CT -segment $q'p$ is crossing free, Case 3 only happens $O(t)$ times before we arrive at Case 1 or Case 2.

Putting the three cases together, we obtain the recurrence

$$f(n) \leq O(t) + f(n'), \text{ with } n' \leq n - 1,$$

which solves to $f(n) = O(tn)$, as claimed. \square

We are now ready to prove the theorem.

Theorem 4.5. Consider a multiset S of n segments with endpoints P , such that P satisfies the property Π that P is partitioned into $P = C \cup T$ where C is in convex position, and $T = \{q, q'\}$ with q outside the convex hull of C and q' inside the convex hull of C . Let t be the sum of the degrees of the points in T , i.e., $t = \deg(q) + \deg(q')$. Let $\mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(n)$ be the number of flips to untangle any multiset of at most n segments with endpoints in convex position, a graph property Γ , and removal choice. There exists a removal strategy R such that any untangle sequence of R for the graph property Γ has length

$$\mathbf{D}_{\Pi, \Gamma}^{\mathbf{R}}(n, t) = O(\deg(q) \deg(q')n + \mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(n)) = O(t^2n + \mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(n)).$$

In particular, we have the following upper bounds for different graph properties:

$$\mathbf{D}_{\Pi}^{\mathbf{R}}(n, t) = O(\deg(q) \deg(q')n + n \log n) = O(t^2n + n \log n) \text{ and}$$

$$\mathbf{D}_{\Pi, \text{TSP}}^{\mathbf{R}}(n), \mathbf{D}_{\Pi, \text{RedBlue}}^{\mathbf{R}}(n) = O(n).$$

Proof. We prove the theorem in the most general case where $\Gamma = \text{Multigraph}$. The untangle sequence is decomposed into four phases.

Phase 1 (CC \times CC). In this phase, we remove all crossings between pairs of CC -segments in S using $\mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(n)$ flips. Throughout all the phases, the invariant that no pair of CC -segments in S crosses is preserved.

Phase 2 (Cq' \times CC). In this phase, we remove all crossings between pairs composed of a CC -segment and a CT -segment incident to q' (the point inside the convex hull of C) using $O(tn)$ flips by Theorem 4.3.

Phase 3 (Cq). At this point, all crossings involve a segment incident to q . In this phase, we deal with all remaining crossings except the crossings involving the segment qq' . Lemma 2.4 allows us to remove the crossings in each CT -segment s incident to q independently, which we do using $O(\deg(q')n)$ flips using Lemma 4.4. As there are $\deg(q)$ CT -segments adjacent to q in S , the total number of flips is $O(\deg(q) \deg(q')n) = O(t^2n)$.

Phase 4 (CC \times TT). At this point, all crossings involve the TT -segment qq' . The endpoints in C that are adjacent to segments with crossings, together with q' , are all in convex position. Hence, the only endpoint not in convex position is q , and we apply Theorem 4.3 using $O(tn)$ flips.

After the $\mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(n)$ flips in Phase 1, the number of flips is dominated by Phase 3 with $O(\deg(q) \deg(q')n) = O(t^2n)$ flips. \square

Notice that, in certain cases (for example in the red-blue case with q, q' having different colors) a flip between two CT -segments never produces two CT -segments. Consequently, Case 3 of the proof of Lemma 4.4 never happens, and the bound in Theorem 4.5 decreases to $O(\mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(n) + tn)$.

4.4 Upper Bound for Two Points Inside a Convex

We prove a similar theorem for two points inside the convex hull of C .

Theorem 4.6. Consider a multiset S of n segments with endpoints P , such that P satisfies the property Π that P is partitioned into $P = C \cup T$ where C is in convex position, and $T = \{q, q'\}$ with q, q' inside the convex hull of C . Let t be the sum of the degrees of the points in T , i.e., $t = \deg(q) + \deg(q')$. Let $\mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(n)$ be the number of flips to untangle any multiset of at most n segments with endpoints in convex position, a graph property Γ , and removal choice. There exists a removal strategy R such that any untangle sequence of R for the graph property Γ has length

$$\mathbf{D}_{\Pi, \Gamma}^{\mathbf{R}}(n, t) = O(tn + \mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(n)).$$

In particular, we have the following upper bounds for different graph properties:

$$\mathbf{D}_{\Pi}^{\mathbf{R}}(n, t) = O(tn + n \log n) \text{ and}$$

$$\mathbf{D}_{\Pi, \text{TSP}}^{\mathbf{R}}(n), \mathbf{D}_{\Pi, \text{RedBlue}}^{\mathbf{R}}(n) = O(n).$$

Proof. The untangle sequence is decomposed in five phases. At the end of each phase, a new type of crossings is removed, and types of crossings removed in the previous phases are not present, even if they may temporarily appear during the phase.

Phase 1 (CT \times CT). In this phase, we remove all crossings between pairs of *CT*-segments in S using $O(\mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(t)) = O(\mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(n))$ flips. We separately solve two convex sub-problems defined by the *CT*-segments in S , one on each side of the line qq' .

Phase 2 (CC \times CC). In this phase, we remove all crossings between pairs of *CC*-segments in S using $O(\mathbf{D}_{\text{Convex}, \Gamma}^{\mathbf{R}}(n))$ flips. As no *CT*-segment has been created, there is still no crossing between a pair of *CT* segments. Throughout, our removal will preserve the invariant that no pair of *CC*-segments in S crosses.

Phase 3 (CT \times non-central CC). We distinguish between a few types of *CC*-segments. The *central* *CC*-segments cross the segment qq' (regardless of qq' being in S or not), while the *non-central* do not. The *peripheral* *CC*-segments cross the line qq' but not the segment qq' , while the *outermost* *CC*-segments do not cross either. In this phase, we remove all crossings between *CT*-segments in S and non-central *CC*-segments in S .

Given a non-central *CC*-segment pp' , let the *out-depth* of the segment pp' be the number of points of C that are contained inside the halfplane bounded by the line pp' and not containing T . Also, let χ be the number of crossings between the non-central *CC*-segments and the *CT*-segments in S . At the end of each step the two following invariants are preserved. (i) No pair of *CC*-segments in S crosses. (ii) No pair of *CT*-segments in S crosses.

At each step, we choose to flip a non-central *CC*-segment pp' of minimum out-depth that crosses a *CT*-segment. We flip pp' with the *CT*-segment $q''p''$ (with $q'' \in \{q, q'\}$) that crosses pp' at the point closest to p (Figure 11(a) and Figure 12(a)). One of the possibly inserted pairs may contain a *CT*-segment s that crosses another *CT*-segment s' , violating the invariant (ii) (Figure 11(b) and Figure 12(b)). If there are multiple such segments s' , then we consider s' to be the segment whose crossing with s is closer to q'' . We flip s and s' and obtain either two *CT*-segments (Figure 11(c) and Figure 12(c)) or a *CC*-segment and the segment qq' (Figure 11(d) and Figure 12(d)). The analysis is divided in two main cases.




Figure 11: Theorem 4.6, Phase 3 when pp' is an outermost segment.

If pp' is an outermost *CC*-segment (see Figure 11), then case analysis shows that the two invariants are preserved and χ decreases.

If pp' is a peripheral *CC*-segment (see Figure 12), then a case analysis shows that the two invariants are preserved and χ has the following behavior. If no *CC*-segment is inserted, then χ decreases. Otherwise a *CC*-segment and a *TT*-segment are inserted and χ may increase by $O(t)$ (Figure 12(d)). Notice that the number of times the *TT*-segment qq' is inserted is $O(t)$, which bounds the total increase by $O(t^2)$.




Figure 12: Theorem 4.6, Phase 3 when pp' is a peripheral segment.

As $\chi = O(tn)$, the total increase is $O(t^2)$, and χ decreases at all but $O(t)$ steps, we have that the number of flips in Phase 3 is $O(tn)$.

Phase 4 ($CT \times CC_{\text{central}}$). At this point, each crossing involves a central CC -segment and either a CT -segment or the TT -segment qq' . In this phase, we remove all crossings between CT -segments and central CC -segments in S , ignoring the TT -segments. This phase ends with crossings only between qq' and central CC -segments in S .




Figure 13: Theorem 4.6, Phase 4. (a) A pair of CT -segments with an ear. (b) A CC -segment and a CT -segment with an ear. (c) Flipping an ear that produces crossing pairs of CT -segments. (d) Flipping an ear that inserts a non-central CC -segment with crossings.

Given four endpoints $q'' \in T$, $p, p'' \in C$, and $x \in C \cup T$, we say that a pair of segments $p''q'', xp \in S$ crossing at a point o contains an *ear* $\widehat{pp''}$ if the interior of the triangle $pp''o$ intersects no segment of S (see Figure 13(a) and 13(b)). Every set of segments with endpoints in $C \cup T$ with $|T| = 2$ that has crossings (not involving the TT -segment) contains an ear (adjacent to the crossing that is farthest from the line qq').

At each *step*, we flip a pair of segments $p''q'', xp$ that contains an ear $\widehat{pp''}$, prioritizing pairs where both segments are CT -segments. Notice that, even though initially we did not have crossing pairs of CT -segments in S , they may be produced in the flip (Figure 13(c)). If the flip inserts a non-central CC -segment which crosses some CT -segments in S (Figure 13(d)), then, we perform the following *while loop*. Assume without loss of generality that qq' is horizontal and s is closer to q' than to q . While there exists a non-central CC -segment s with crossings, we flip s with the CT -segment s' crossing s that comes first according to the following order. As a first criterion, a segment incident to q comes before a segment incident to q' . As a second tie-breaking criterion, a segment whose crossing point with s that is farther from the line qq' comes before one that is closer.

Let $\chi = O(tn)$ be the number of crossings between central CC -segments and CT -segments in S plus the number of crossings between CT -segments in S . A case analysis shows that the value of χ decreases at each step. If no non-central CC -segment is inserted, then the corresponding step consists of a single flip. As χ decreases, there are $O(tn)$ steps that do not insert a non-central CC -segment.

However, if a non-central CC -segment is inserted, at the end of the step we inserted a CC -segment that is uncrossable (see Lemma 2.4). As the number of CC -segments in S is $O(n)$, we have that the number of times the while loop is executed is $O(n)$. Since each execution of the while loop performs $O(t)$ flips, we have a total of $O(tn)$ flips in this phase.

Phase 5 ($TT \times CC_{\text{central}}$). In this phase, we remove all crossings left, which are between the possibly multiple copies of the TT -segment qq' and central CC -segments in S . The endpoints of the segments with crossings are in convex position and all other endpoints are outside their convex hull. Hence, by Lemma 2.4, it is possible to obtain a crossing-free multigraph using $O(\mathbf{D}_{\text{Convex},\Gamma}^{\text{R}}(n))$ flips. \square

4.5 Upper Bound for Two Points Outside a Convex

In this section, we consider the case of two endpoints q, q' outside the convex hull of the remaining endpoints C , which are in convex position. We prove a theorem with a bound that is quadratic in $t = \deg(q) + \deg(q')$ but linear or near-linear in n .

Theorem 4.7. *Consider a multiset S of n segments with endpoints P , such that P satisfies the property Π that P is partitioned into $P = C \cup T$ where C is in convex position, and $T = \{q, q'\}$ with q, q' outside the convex hull of C . Let t be the sum of the degrees of the points in T , i.e., $t = \deg(q) + \deg(q')$. Let $\mathbf{D}_{\text{Convex},\Gamma}^{\text{R}}(n)$ be the number of flips to untangle any multiset of at most n segments with endpoints in convex position, a graph property Γ , and removal choice. There exists a removal strategy R such that any untangle sequence of R for the graph property Γ has length at most*

$$\mathbf{D}_{\Pi,\Gamma}^{\text{R}}(n, t) = O(\deg(q) \deg(q') n) + \mathbf{D}_{\text{Convex},\Gamma}^{\text{R}}(n) = O(t^2 n + \mathbf{D}_{\text{Convex},\Gamma}^{\text{R}}(n)).$$

In particular, we have the following upper bounds for different graph properties:

$$\mathbf{D}_{\Pi}^{\text{R}}(n, t) = O(t^2 n + n \log n) \text{ and}$$

$$\mathbf{D}_{\Pi,\text{TSP}}^{\text{R}}(n), \mathbf{D}_{\Pi,\text{RedBlue}}^{\text{R}}(n) = O(n).$$




Figure 14: Illustration of the four case of relative position of the points q, q' and the convex hull of C for the proof of Theorem 4.7. Without loss of generality, the case (d) is not considered.

Without loss of generality, we assume that q' is not in the interior of the convex hull of $C \cup \{q\}$ (Figure 14). The strategy starts with the following preprocessing phase. We *delete* from S all segments incident to q . We then untangle S using Theorem 4.3, which uses at most $(n-1) \deg(q') + \mathbf{D}_{\text{Convex},\Gamma}^{\text{R}}(n) = O(n \deg(q')) + \mathbf{D}_{\text{Convex},\Gamma}^{\text{R}}(n)$ flips.

Next, we *add back* to S all segments incident to q , one by one, in any order. When adding back a TT -segment (Figure 15(a)), we call Routine TT , which will untangle S using $O(n)$ flips (Figure 15(d)). When adding back a CT -segment (Figure 16(a) and Figure 17(a)), we call Routine CT , which will untangle S using $O(n \deg(q'))$ flips. Since the number of segments that we add back is $\deg(q)$, the bound stated in Theorem 4.7 holds. Next, we describe Routine TT , followed by Routine CT .




Figure 15: Illustration for Routine TT in the proof of Theorem 4.7. (a) A TT -segment (in bold) is added back to S . (b) The first flip (the flipped segments are in bold, the inserted segments are dashed in red). The segments of S that have been removed from S' are grayed. The splitting line is highlighted. (c) The remaining flips. (d) S is crossing free at the end of Routine TT .

Routine TT . If qq' has no crossing, then we are done as S is crossing free. Otherwise, we know that there is a single copy of qq' , because $S \setminus \{qq'\}$ is crossing free. We start by flipping qq' with an arbitrary segment pp' that crosses qq' . Assume without loss of generality that the flip inserted pq and $p'q'$ (Figure 15(b)).

We now construct a splitting partition of S in order to apply Lemma 2.4. Let S' be the set of segments that have not been assigned a partition yet. Initially $S' = S$. We say that a segment is S' -uncrossable if it is not crossed by any segment defined by two endpoints of segments in S' . While there is an S' -uncrossable segment $u \in S'$, we create a new singleton partition containing u , removing u from S' . We claim that at the end of this loop, the line pp' is crossing free (Figure 15(b)). We then create two separate partitions S_1, S_2 with the segments of S' on each side of pp' , as well as a partition with the copies of the segment pp' that are still in S' (if any). We are now ready to apply Lemma 2.4. The only partitions that may have crossings are S_1, S_2 . Each of S_1, S_2 have all crossings in the single CT -segment. Hence, we apply Theorem 4.3 with $t = 1$ to untangle each of S_1, S_2 using $O(n)$ flips, and consequently untangling S using $O(n)$ flips (Figure 15(c)).

We say that a CT -segment is a CT_{in} -segment if it intersects the interior of the convex hull of C and a CT_{out} -segment otherwise. To prove the claim, first note that, since the segment qq' intersects the convex hull of C , all the CT_{out} -segments are initially (i.e., when $S' = S$) uncrossable, and thereby removed first from S' . Second, after all the CT_{out} -segments have been removed from S' , consider the CT_{in} -segment s'' of S' whose endpoint in C is the farthest away from the line qq' . The line ℓ'' containing s'' is crossing free (Figure 15(b)). Thus, by Lemma 2.4, we partition S' in three: the segments on each side of the line ℓ'' and the segments on the line ℓ'' itself. The segments on the side of ℓ'' which does not contain the open segment qq' are crossing free, and are thus removed from S' . The claim follows, since at the end of the loop, all the CT -segments have been removed from S' .

Routine CT . This routine untangles S in the case where the segment s that we add back to S is a CT -segment $s = pq$ (Figure 16(a) and Figure 17(a)). This routine consists of the following while loop and calls a subroutine called Restore-Loop-Invariant to maintain the invariant that removing a single CT -segment from S would remove all the crossings.

While the segment s crosses some segment in S :

1. Let s' be the segment of S that crosses s the farthest away from q .
2. Flip s with s' (see Figure 16(b,c) and Figure 17(b,c,d) for five different cases of this flip).
3. If s' is not a CC -segment, then call Restore-Loop-Invariant.




Figure 16: Illustrations for Routine CT in the proof of Theorem 4.7. (a) A CT -segment (in bold) is added back to S . (b) The case where s' is a CC -segment. (c) The case where s' is a CT -segment. The existence of the two grayed segments is assumed for a contradiction. (d) Illustration of the subroutine Restore-Loop-Invariant.

4. Set s to be the segment incident to q that was inserted by the flip of s with s' .

Restore-Loop-Invariant. Let o be the crossing point of $s = qp$ and $s' = p'q'$.

1. If the inserted pair is qp', pq' (Figure 16(c) and Figure 17(b)):
 - (a) If pq' has crossings, then untangle the convex sector defined by the rays oq' and op , i.e., the rays with origin o and directed by the vectors $q' - o$ and $p - o$ (Figure 16(d)). We will show that these rays are crossing free. We then apply Theorem 4.3 in the case where the CC -segments are already crossing free and $t = 1$ (as the other segments incident to q' are already crossing free), using $O(n)$ flips.
2. If the inserted pair is qq', pp' (we will show that the following cases are mutually exclusive):
 - (a) If qq' has crossings (Figure 17(c)), then call Routine TT (at most n flips).
 - (b) If pp' has crossings (Figure 17(d)), then pp' must cross another CC -segment. While there is a crossing between two CC -segments, we flip them, in any order. This procedure does not create crossings involving CT -segments or TT -segments, hence the number of crossings strictly decreases. Consequently, the number of flips performed is at most the number of crossings of the segment pp' when the routine was called, which is $O(n)$.

Next, we show that the rays oq' and op in step 1a are crossing free, which ensures that flips performed in this convex sector do not insert segments that cross the remaining segments. The proof breaks the two rays into two parts each.

Since the segment $s' = p'q'$ is crossing free, then the segment oq' is crossing free. Since s' is the segment that crosses s the farthest away from q , then the segment op is crossing free.

At this point, we know that all the segments of S crossing the segment pq' are CC -segments. Let o_1 be the intersection point between pq' and an arbitrary CC -segment. The remaining part of the ray oq' (that is, the ray oq' deprived of the segment oq') is crossing free because, by definition, q' is not in the interior of the convex hull of $C \cup \{q\}$.

It only remains to prove that the ray op deprived of the segment op is also crossing free. For the sake of a contradiction, we assume that this is not the case. If a CC -segment crosses the ray op at a point o_2 , then p is in the interior of the triangle o_1o_2p' (Figure 16(c)). If a CT -segment p_2q' crosses the ray op , then p is in the interior of the triangle o_1p_2p' (Figure 16(c)). Both cases contradicts the fact that $p \in C$.




Figure 17: Illustrations for Routine CT in the proof of Theorem 4.7. The existence of the grayed segments is assumed for a contradiction. (a) A CT -segment (in bold) is added back to S (in an alternative context compared to Figure 16). (b) Illustration for the proof of Lemma 4.8; in the case where s' is a CT -segment and the inserted segment $p'q$ crosses a CC -segment at o_4 . (c) & (d) The two cases where the inserted segments are qq' and pp' , whether the segment qq' intersects the convex hull of C .

We now show that the two cases of step 2 are mutually exclusive.

For the segment pp' to have crossings, at least one point of P has to be in the interior of the triangle $pp'o$. We examine the following two cases, which are mutually exclusive by definition, and show that the two cases of step 2 respectively imply case 2a' and case 2b'.

Case 2a': In this case, there exists a point o_3 of the TT -segment qq' that is in the interior of the convex hull of C (Figure 17(c)). By convexity, the interior of the triangle $pp'o_3$ does not contain any point of P . The triangle $pp'o_3$ contains the triangle $pp'o$, and the segments op and op' are crossing free. Therefore, the CC -segment pp' is crossing free.

Case 2b': In this case, the TT -segment qq' does not intersect the interior of the convex hull of C (Figure 17(d)). Consequently, the TT -segment qq' is crossing free. In this case, all the crossings of S involve the CC -segment pp' and other CC -segments.

Number of flips. We now analyze the number of flips performed by Routine CT , which includes the flips in the subroutine Restore-Loop-Invariant.

Lemma 4.8. *Each call to Routine CT performs $O(n \deg(q'))$ flips.*

Proof. We first limit our attention to the sequence Q of flips between s and s' at line 2 of Routine CT . There are two types of flips performed, either s' is a CC -segment or a CT -segment. In the following, we show that:

1. The number of consecutive flips in Q where s' is a CC -segment is at most $n - 1$.
2. The total number of flips in Q where s' is a CT -segment is at most $\deg(q')$.

Lemma 4.8 follows from 1 and 2, and from the fact that each call to Restore-Loop-Invariant performs $O(n)$ flips. Indeed, the flips performed by Routine CT between two flips where s' is a CT -segment are either consecutive flips where s' is a CC -segment or flips performed by a call to Restore-Loop-Invariant.

To prove 1, notice that the proof of Lemma 4.2 applied to the CC -segments of S and s ensures that the number of consecutive flips where s' is a CC -segment is at most $n - 1$.

We now prove 2. We apply Lemma 2.4 at the end of a maximal subsequence Q' of Q consisting of flips where s' is a CT -segment. Let k be the length of Q' . We first observe that k is at most the number of CT -segments crossing s just before the first flip of this sequence. This observation relies on the correctness of Restore-Loop-Invariant and on the following simple geometric fact.

Given a triangle $qq'p'$ and a point o on the segment $p'q'$, any segment incident to q' and crossing the segment $p'q$ also crosses the segment oq .

We now consider the configuration just after the last flip of Q' . If $p'q$ is crossing free, then S is in fact crossing free and Routine CT ends. Otherwise, we have a splitting partition for Lemma 2.4, consisting of at most three sets of segments, one for each side of the line $q'p'$, and one for the segments on the line, if any. Next, we show that the line $q'p'$ is crossing free (see Figure 17(b) for an illustration of this splitting partition and of the notations introduced in the proof).

Let o_4 be the crossing point of $p'q$ and the CC -segment crossing $p'q$ the farthest away from q . For the sake of a contradiction, we assume that the line $q'p'$ is not crossing free. If the line $q'p'$ crosses a CC -segment at a point o_5 , then p' is in the interior of the triangle o_4o_5p (Figure 17(b)). If the line $q'p'$ crosses a CT -segment qp_5 , then p' is in the interior of the triangle o_4p_5p (Figure 17(b)). Both cases contradict the fact that $p' \in C$.

The partition containing the point p is crossing free and we claim that it contains at least k segments incident to q' and it . Thus, 2 follows. To prove the claim, it is enough to observe that, at each flip of Q where s' is a CT -segment, s always rotates in the same direction. More formally, the sign of the measure between $-\pi$ and π of the oriented angle from the vector $q - p$ to the vector $q - p'$ is always the same. \square

4.6 Upper Bound for Two Points Inside or Outside a Convex

The most important results of Theorems 4.5, 4.6, and 4.7 are summarized in the following theorem.

Theorem 4.9. *Consider a multiset S of n segments with endpoints P , such that P satisfies the property Π that P is partitioned into $P = C \cup T$ where C is in convex position, and $T = \{q, q'\}$. Let t be the sum of the degrees of the points in T , i.e., $t = \deg(q) + \deg(q')$. Let $\mathbf{D}_{\text{Convex}, \Gamma}^{\text{R}}(n)$ be the number of flips to untangle any multiset of at most n segments with endpoints in convex position, a graph property Γ , and removal choice. There exists a removal strategy R such that any untangle sequence of R for the graph property Γ has length*

$$\mathbf{D}_{\Pi, \Gamma}^{\text{R}}(n, t) = O(t^2n + \mathbf{D}_{\text{Convex}, \Gamma}^{\text{R}}(n)).$$

In particular, we have the following upper bounds for different graph properties:

$$\mathbf{D}_{\Pi}^{\text{R}}(n, t) = O(t^2n + n \log n) \text{ and}$$

$$\mathbf{D}_{\Pi, \text{TSP}}^{\text{R}}(n), \mathbf{D}_{\Pi, \text{RedBlue}}^{\text{R}}(n) = O(n).$$

5 Untangling with Insertion Choice

In this section, we devise strategies for insertion choice to untangle multisets of segments, therefore providing upper bounds on several versions of \mathbf{D}^{I} . Recall that such insertion strategies *do not* choose which pair of crossing segments is removed, but only which pair of segments with the same endpoints is subsequently inserted. We start with a tight bound in **Convex** version (where the point set is in convex position), followed by the version where the point set $P = C \cup T$ has some points C in convex position and the other points T outside the convex hull of C and separated from C by two parallel lines.

5.1 Upper Bound for Convex Position

Let $P = C = \{p_1, \dots, p_{|C|}\}$ be a set of points in convex position sorted in counterclockwise order along the convex hull boundary (Figure 18(a)). Given a segment $p_a p_b$, we define the *depth* $\delta(p_a p_b) = |b - a|$.¹ We use the depth to prove the following theorem.




Figure 18: (a) A multigraph (C, S) with $|C| = 14$ points in convex position and $n = 9$ segments. (b) Insertion choice for Case 1 and 2 of the proof of Theorem 5.1. (c) Insertion choice for Case 3.

Theorem 5.1. *Consider a multiset S of n segments with endpoints $P = C$ in convex position. There exists an insertion strategy I such that any untangle sequence of I has length*

$$\mathbf{D}_{\text{Convex}}^I(n) = O(n \log |P|) = O(n \log n).$$

Proof. Let the potential function

$$\varpi(S) = \prod_{s \in S} \delta(s).$$

As $\delta(s) \in \{1, \dots, |C| - 1\}$, we have that $\varpi(S)$ is integer, positive, and at most $|C|^n$. Next, we show that for any flipped pair of segments $p_a p_b, p_c p_d$ there exists an insertion choice that multiplies $\varpi(S)$ by a factor of at most $3/4$, and the theorem follows.

Consider a flip of a segment $p_a p_b$ with a segment $p_c p_d$ and assume without loss of generality that $a < c < b < d$. The contribution of the pair of segments $p_a p_b, p_c p_d$ to the potential $\varpi(S)$ is the factor $f = \delta(p_a p_b) \delta(p_c p_d)$. Let f' be the factor corresponding to the pair of inserted segments.

Case 1: If $\delta(p_a p_c) \leq \delta(p_c p_b)$, then we insert the segments $p_a p_c$ and $p_b p_d$ and we get $f' = \delta(p_a p_c) \delta(p_b p_d)$ (Figure 18(b)). We notice $\delta(p_a p_b) = \delta(p_a p_c) + \delta(p_c p_b)$. It follows $\delta(p_a p_c) \leq \delta(p_a p_b)/2$ and we have $\delta(p_b p_d) \leq \delta(p_c p_d)$ and then $f' \leq f/2$.

Case 2: If $\delta(p_b p_d) \leq \delta(p_c p_b)$, then we insert the same segments $p_a p_c$ and $p_b p_d$ as previously. We have $\delta(p_a p_c) \leq \delta(p_a p_b)$ and $\delta(p_b p_d) \leq \delta(p_c p_d)/2$, which gives $f' \leq f/2$.

Case 3: If (i) $\delta(p_a p_c) > \delta(p_c p_b)$ and (ii) $\delta(p_b p_d) > \delta(p_c p_b)$, then we insert the segments $p_a p_d$ and $p_c p_b$ (Figure 18(c)). The contribution of the new pair of segments is $f' = \delta(p_a p_d) \delta(p_c p_b)$. We introduce the coefficients $x = \frac{\delta(p_a p_c)}{\delta(p_c p_b)}$ and $y = \frac{\delta(p_b p_d)}{\delta(p_c p_b)}$ so that $\delta(p_a p_c) = x \delta(p_c p_b)$ and $\delta(p_b p_d) = y \delta(p_c p_b)$. It follows that $\delta(p_a p_b) = (1 + x) \delta(p_c p_b)$, $\delta(p_c p_d) = (1 + y) \delta(p_c p_b)$ and $\delta(p_a p_d) = (1 + x + y) \delta(p_c p_b)$. The ratio f'/f is equal to a function $g(x, y) = \frac{1+x+y}{(1+x)(1+y)}$. Due to (i) and (ii), we have that $x \geq 1$ and $y \geq 1$. In other words, we can upper bound the ratio f'/f by the maximum of the function $g(x, y)$ with $x, y \geq 1$. It is easy to show that the function $g(x, y)$ is decreasing with both x and y . Then its maximum is obtained for $x = y = 1$ and it is equal to $3/4$, showing that $f' \leq 3f/4$. \square

¹This definition resembles but is not exactly the same as the depth used in [10]. It is also similar to the crossing depth defined in Section 4.

5.2 Upper Bound for Points Separated by Two Parallel Lines

In this section, we prove the following theorem, which is a generalization of Theorem 5.1. We extend our standard general position assumptions to also exclude pairs of endpoints with the same y -coordinate.




Figure 19: (a) Statement of Theorem 5.2. (b) Some insertion choices in the proof of Theorem 5.2.

Theorem 5.2. *Consider a multiset S of n segments with endpoints P , such that P satisfies the property Π that P is partitioned into $P = C \cup T_1 \cup T_2$ where C is in convex position and there exist two horizontal lines ℓ_1, ℓ_2 , with T_1 above ℓ_1 above C above ℓ_2 above T_2 . Let t be the sum of the degrees of the points in $T = T_1 \cup T_2$. There exists an insertion strategy I such that any untangle sequence of I has length*

$$D_{\Pi}^I(n, t) = O(t|P| \log |C| + n \log |C|) = O(tn \log n).$$

Proof. We start by describing the insertion choice for flips involving at least one point in T . Let $p_1, \dots, p_{|P|}$ be the points P sorted vertically from top to bottom. Consider a flip involving the points p_a, p_b, p_c, p_d with $a < b < c < d$. The insertion choice is to create the segments $p_a p_b$ and $p_c p_d$. See Figure 19(b). As in [12] (specifically, in the proof of the upper bound of Theorem 2), we define the potential η of a segment $p_i p_j$ as

$$\eta(p_i p_j) = |i - j|.$$

Notice that η is an integer between 1 and $|P| - 1$. We define $\eta_T(S)$ as the sum of $\eta(p_i p_j)$ for $p_i p_j \in S$ with p_i or p_j in T . Notice that $0 < \eta_T(S) < t|P|$. It is easy to verify that any flip involving a point in T decreases $\eta_T(S)$ and other flips do not change $\eta_T(S)$. Hence, the number of flips involving at least one point in T is $O(t|P|)$.

For the flips involving only points of C , we use the same choice as in the proof of Theorem 5.1. The potential function

$$\varpi(S) = \prod_{p_i p_j \in S : p_i \in C \text{ and } p_j \in C} \delta(p_i p_j)$$

is at most $|C|^n$ and decreases by a factor of at most $3/4$ at every flip that involves only points of C .

However, $\varpi(S)$ may increase by a factor of $O(|C|^2)$ when performing a flip that involves a point in T . As such flips only happen $O(t|P|)$ times, the total increase is at most a factor of $|C|^{O(t|P|)}$.

Concluding, the number of flips involving only points in C is at most

$$\log_{4/3} \left(|C|^{O(n)} |C|^{O(t|P|)} \right) = O(n \log |C| + t|P| \log |C|). \quad \square$$

6 Untangling with Both Choices

In this section, we devise strategies for insertion and removal choices to untangle a multiset of segments, therefore providing upper bounds on \mathbf{D}^{RI} with some endpoints outside (but not inside) a convex polygon. Recall that such strategies choose which pair of crossing segments is removed *and* which pair of segments with the same endpoints is subsequently inserted.

Throughout this section, we assume that the point set P is partitioned into $P = C \cup T$ where C is in convex position and the points T lie outside the convex hull of C . We start with the case where T is separated by two parallel lines from C . Afterwards, we prove an important lemma and apply it to untangle a matching.

6.1 Upper Bound for Points Separated by Two Parallel Lines

In this section, we prove an upper bound when T is separated from C by two parallel lines. In this version, our bound of $O(n + t|P|)$ interpolates the tight convex bound of $O(n)$ from [10, 20] and the $O(t|P|)$ bound from [12] for t arbitrary segments. We extend our standard general position assumptions to also exclude pairs of endpoints with the same y -coordinate.

Theorem 6.1. *Consider a multiset S of n segments with endpoints P , such that P satisfies the following property Π . The endpoints P are partitioned into $P = C \cup T_1 \cup T_2$ where C is in convex position and there exist two horizontal lines ℓ_1, ℓ_2 , with T_1 above ℓ_1 above C above ℓ_2 above T_2 . Let t be the sum of the degrees of the points in $T = T_1 \cup T_2$.*

There exists a removal and insertion strategy RI such that any untangle sequence of RI has length

$$\mathbf{D}_{\Pi}^{\text{RI}}(n, t) = O(n + t|P|) = O(tn).$$

Proof. The algorithm runs in two phases.

Phase 1. We use removal choice to perform the flips involving a point in T . At the end of the first phase, there can only be crossings among segments with all endpoints in C . The insertion choice for the first phase is the following. Let $p_1, \dots, p_{|P|}$ be the points P sorted vertically from top to bottom. Consider a flip involving the points p_a, p_b, p_c, p_d with $a < b < c < d$. The insertion choice is to create the segments $p_a p_b$ and $p_c p_d$. As in the proofs of Theorem 5.1 and of Theorem 5.2, we define the potential δ of a segment $p_i p_j$ as $\delta(p_i p_j) = |i - j|$.² Notice that δ is an integer from 1 to $|P| - 1$. We define $\delta(S)$ as the sum of $\delta(p_i p_j)$ for $p_i p_j \in S$ with p_i or p_j in T . Notice that $0 < \delta(S) < t|P|$. It is easy to verify that any flip involving a point in T decreases $\delta(S)$. Hence, the number of flips in Phase 1 is $O(t|P|)$.

Phase 2. Since T is outside the convex hull of C , flips between segments with all endpoints in C cannot create crossings with the other segments, which are guaranteed to be crossing free at this point. Hence, it suffices to run the algorithm to untangle a convex set with removal and insertion choice from [10]³, which performs $O(n)$ flips. \square

6.2 Upper Bound for Matchings with Points Outside a Convex

In this section, we consider the case of endpoints $C \cup T$ with C in convex position and T outside C . The result only apply to matching because Lemma 6.2 is false when multiple copies of a segment are allowed. We start by proving Lemma 6.2 and then apply it to prove Theorem 6.4.

²This definition resembles but is not the same as the depth used in [10].

³The algorithm in [10] is originally proven for convex matchings, but we use here a straightforward generalization to multigraphs available in [40], Theorem 3.2.13.

6.2.1 Liberating a Line

In this section, we prove the following key lemma, which we use in the following section. The lemma only applies to matchings and it is easy to find a counter-example for multisets (S consisting of n copies of a single segment that crosses pq).

Lemma 6.2. *Consider a set S of n segments with endpoints C in convex position and a segment qq' intersecting the interior of the convex hull of C such that $(C \cup \{q, q'\}, S \cup \{qq'\})$ forms a matching.*

There exists a flip sequence starting at $S \cup \{qq'\}$ of length $O(n)$ which ends with a set of segments that do not cross the line qq' (the line qq' splits the final set of segments).

Proof. For each flip performed in the subroutine described hereafter, at least one of the inserted segments does not cross the line qq' and is removed from S (see Figure 20).




Figure 20: An untangle sequence of the subroutine to liberate the line qq' (with $n = 4$).

Preprocessing. First, we remove from S the segments that do not intersect the line qq' , as they are irrelevant. Second, anytime two segments in S cross, we flip them choosing to insert the pair of segments not crossing the line qq' . One such flip removes two segments from S . Let p_1p_2 (respectively $p_{2n-1}p_{2n}$) be the segment in S whose intersection point with qq' is the closest from q (respectively q'). Without loss of generality, assume that the points p_1 and p_{2n-1} are on the same side of the line qq' .

First flip. Lemma 6.3 applied to the segment qq' and the triangle $p_1p_2p_{2n-1}$ shows that at least one of the segments among $qp_{2n-1}, q'p_1, q'p_2$ intersects all the segments of S . Without loss of generality, assume that qp_{2n-1} is such a segment, i.e., that qp_{2n-1} crosses all segments of $S \setminus \{p_{2n-1}p_{2n}\}$. We choose to remove the segments qq' and $p_{2n-1}p_{2n}$, and we choose to insert the segments qp_{2n-1} and $q'p_{2n}$. As the segment $q'p_{2n}$ does not cross the line qq' , we remove it from S .

Second flip. We choose to flip the segments qp_{2n-1} and p_1p_2 . If n is odd, we choose to insert the pair of segments qp_1, p_2p_{2n-1} . If n is even, we insert the segments qp_2, p_1p_{2n-1} .

By convexity, one of the inserted segment (the one with endpoints in C) crosses all other $n - 2$ segments. The other inserted segment (the one with q as one of its endpoints) does not cross the line qq' , so we remove it from S . Note that the condition on the parity of n is there only to ensure that the last segment $p_{2n-3}p_{2n-2}$ is dealt with at the last flip.

Remaining flips. We describe the third flip. The remaining flips are performed similarly. Let s be the previously inserted segment. Let p_3p_4 be the segment in S whose intersection point with qq' is the closest from q . Without loss of generality, assume that p_3 is on the same side of the line qq' as p_1 and p_{2n-1} .

We choose to flip s with p_3p_4 . If $s = p_2p_{2n-1}$, we choose to insert the pair of segments p_2p_4, p_3p_{2n-1} . If $s = p_1p_{2n-1}$, we choose to insert the pair of segments p_1p_3, p_4p_{2n-1} .

By convexity, one inserted segment (the one with p_{2n-1} as an endpoint) crosses all other $n - 3$ segments. The other inserted segment does not cross the line qq' , so we remove it from S . Note that the insertion choice described is the only viable one, as the alternative would insert a crossing-free segment crossing the line qq' that cannot be removed. \square

Auxiliary Lemma of Section 6.2.1. In this section, we prove Lemma 6.3 used in the proof of Lemma 6.2.

Recall that, in the proof of Lemma 6.2, we have a convex quadrilateral $p_1p_2p_{2n}p_{2n-1}$ and a segment qq' crossing the segments p_1p_2 and $p_{2n}p_{2n-1}$ in this order when drawn from q to q' , and we invoke Lemma 6.3 to show that at least one of the segments among $qp_{2n-1}, q'p_1, q'p_2$ intersects all the segments of S . Before proving Lemma 6.3, we detail how to apply it to this context.

Lemma 6.3 applied to the segment qq' and the triangle $p_1p_2p_{2n-1}$ asserts that at least one of the following pairs of segments cross: qp_{2n-1}, p_1p_2 , or $q'p_1, p_2p_{2n-1}$, or $q'p_2, p_1p_{2n-1}$. If the segments qp_{2n-1}, p_1p_2 cross, then we are done. If the segments $q'p_1, p_2p_{2n-1}$ cross, then the segments $q'p_1, p_{2n}p_{2n-1}$ also cross and we are done. If the segments $q'p_2, p_1p_{2n-1}$ cross, then the segments $q'p_2, p_{2n}p_{2n-1}$ also cross and we are done.

Next, we state and prove Lemma 6.3.

Lemma 6.3. *For any triangle $p_1p_2p_3$, for any segment qq' intersecting the interior of the triangle $p_1p_2p_3$, there exists a segment $s \in \{qp_1, qp_2, qp_3, q'p_1, q'p_2, q'p_3\}$ that intersects the interior of the triangle $p_1p_2p_3$.*

Proof. If all p_1, p_2, p_3, q, q' are in convex position, then q and the point among p_1, p_2, p_3 that is not adjacent to q on the convex hull boundary define the segment s . Otherwise, since q, q' are not adjacent on the convex hull boundary, assume without loss of generality that p_1 is not a convex hull vertex and q, p_2, q', p_3 are the convex hull vertices in order. Then, either the segment p_1q or the segment p_1q' intersects the segment p_2p_3 . \square

6.2.2 Proof of the Upper Bound

We are now ready to prove the following theorem, which only applies to matchings because it uses Lemma 6.2.

Theorem 6.4. *Consider a set S of n segments with endpoints P partitioned into $P = C \cup T$ where C is in convex position and T is outside the convex hull of C and such that (P, S) defines a matching. Let $t = |T|$.*

There exists a removal and insertion strategy RI such that any untangle sequence of RI has length

$$\mathbf{D}_{\text{Matching}}^{\text{RI}}(n, t) = O(t^3n).$$

Proof. Throughout this proof, we partition the TT -segments into two types: TT_{in} -segment if it intersects the interior of the convex hull of C and TT_{out} -segment otherwise. We define the potential $\Lambda_\ell(S)$ of a line ℓ as the number of segments of S crossing ℓ .

TT -segments. At any time during the untangle procedure, if there is a TT_{in} -segment s that crosses more than t segments, we apply Lemma 6.2 to liberate s from every CC -segment using $O(n)$ flips. Let ℓ be the line containing s . Since Λ_ℓ cannot increase (by an easy observation that first appears in [42]), $\Lambda_\ell < t$ after Lemma 6.2, and there are $O(t^2)$ different TT_{in} -segments in S , it follows that Lemma 6.2 is applied $O(t^2)$ times, performing a total $O(t^2n)$ flips. As the number

of times s is inserted and removed differ by at most 1 and Λ_ℓ decreases at each flip that removes s , it follows that s participates in $O(t)$ flips. As there are $O(t^2)$ different TT_{in} -segments in S , the total number of flips involving TT_{in} -segments is $O(t^3)$.

We define a set L of $O(t)$ lines as follows. For each point $q \in T$, we have two lines $\ell_1, \ell_2 \in L$ that are the two tangents of the convex hull of C that pass through q . As the lines $\ell \in L$ do not intersect the interior of the convex hull of C , the potential $\Lambda_\ell = O(t)$. When flipping a TT_{out} -segment q_1q_2 with another segment q_3p with $q_3 \in T$ (p may be in T or in C), we make the insertion choice of creating a TT_{out} -segment q_1q_3 such that there exists a line $\ell \in L$ whose potential Λ_ℓ decreases. It is easy to verify that ℓ always exist (see Lemma 6.5 and Lemma 6.6). Hence, the number of flips involving TT_{out} -segments is $O(t^2)$ and the number of flips involving TT -segments in general is $O(t^3)$.

All except pairs of CC -segments. We keep flipping segments that are not both CC -segments with the following insertion choices. Whenever we flip two CT -segments, we make the insertion choice of creating a TT -segment. Hence, as the number of flips involving TT -segments is $O(t^3)$, so is the number of flips of two CT -segments.

Whenever we flip a CT -segment p_1q with $q \in T$ and a CC -segment p_3p_4 , we make the following insertion choice. Let $v(q)$ be a vector such that the dot product $v(q) \cdot q < v(q) \cdot p$ for all $p \in C$, that is, v is orthogonal to a line ℓ separating q from C and v is pointing towards C . We define the potential $\rho(p_iq)$ of a segment with $p_i \in C$ and $q \in T$ as the number of points $p \in C$ such that $v(q) \cdot p < v(q) \cdot p_i$, that is the number of points in C before p_i in direction v . We choose to insert the segment p_iq that minimizes $\rho(p_iq)$ for $i \in \{3, 4\}$. Let $\rho(S)$ be the sum of $\rho(p_iq)$ for all CT -segments p_iq in S . It is easy to see that $\rho(S)$ is $O(t|C|)$ and decreases at each flip involving a CT -segment (not counting the flips inside Lemma 6.2).

There are two situation in which $\rho(S)$ may increase. One is when Lemma 6.2 is applied, which happens $O(t^2)$ times. Another one is when a TT -segment and a CC -segment flip, creating two CT -segments, which happens $O(t^3)$ times. At each of these two situations, $\rho(S)$ increases by $O(|C|)$. Consequently, the number of flips between a CT -segment and a CC -segment is $O(t^3|C|) = O(t^3n)$.

CC -segments. By removal choice, we choose to flip the pairs of CC -segments last (except for the ones flipped in Lemma 6.2). As T is outside the convex hull of C , flipping two CC -segments does not create crossings with other segments (by Lemma 2.4). Hence, we apply Theorem 5 from [10] to untangle the remaining segments using $O(n)$ flips. \square

Auxiliary Lemmas of Section 6.2 In this section, we prove Lemma 6.6 and Lemma 6.5 used in the proof of Theorem 6.4.

Recall that, in the proof of Theorem 6.4, we define a set L of lines as follows. For each point $q \in T$, we have two lines $\ell_1, \ell_2 \in L$ that are the two tangents of the convex hull of C that pass through q . When flipping a TT_{out} -segment q_1q_2 with another segment q_3p with $q_3 \in T$ (p may be in T or in C), we make the insertion choice of creating a TT_{out} -segment q_1q_3 such that there exists a line $\ell \in L$ whose potential Λ_ℓ decreases. We invoke Lemma 6.5 and Lemma 6.6 to show that such a line ℓ always exist.

Indeed, by Lemma 6.5, it is enough to show that there exists a line $\ell \in L$ containing one of the points q_1, q_2, q_3 that crosses one of the segments q_1q_2 or q_3p . This is precisely what Lemma 6.6 shows.

Next, we state prove Lemma 6.5 and Lemma 6.6.

Lemma 6.5. *Consider two crossing segments p_1p_2, p_3p_4 and a line ℓ containing p_1 and crossing p_3p_4 . Then, one of the two pairs of segments p_1p_3, p_2p_4 or p_1p_4, p_2p_3 does not cross ℓ . In other*

words, there exists an insertion choice for a flip removing p_1p_2, p_3p_4 such that the number of segments crossing ℓ decreases.

Proof. Straightforward. □

Lemma 6.6. *Consider a closed convex body B and two crossing segments q_1q_3, q_2q_4 whose endpoints q_1, q_2, q_3 are not in B , and whose endpoint q_4 is not in the interior of B . If the segment q_1q_3 does not intersect the interior of B , then at least one of the six lines tangent to B and containing one of the endpoints q_1, q_2, q_3 is crossing one of the segments q_1q_3, q_2q_4 (Figure 21(a)). (General position is assumed, meaning that the aforementioned six lines are distinct, i.e., each line does not contain two of the points q_1, q_2, q_3, q_4 .)*




Figure 21: (a) In the statement of Lemma 6.6, we assert the existence of points, circled in the figure, which are the intersection of a line tangent to B and containing one of the points q_1, q_2, q_3 . (b) In the proof of Lemma 6.6 by contraposition, we exhibit a point, circled in the figure, showing that B intersects one of the segment q_1q_3 .

Proof. For all $i \in \{1, 2, 3\}$, let ℓ_i and ℓ'_i be the two lines containing q_i and tangent to B (Figure 21(b)). By contraposition, we assume that none of the six lines $\ell_1, \ell'_1, \ell_2, \ell'_2, \ell_3, \ell'_3$ crosses one of the segments q_1q_3, q_2q_4 . In other words, we assume that the six lines are tangent to the convex quadrilateral $q_1q_2q_3q_4$. It is well known that, if $m \geq 5$, then any arrangement of m lines or more admits at most one face with m edges (see [26] for example). Therefore, B is contained in the same face of the arrangement of the six lines as the quadrilateral $q_1q_2q_3q_4$. Let o_1 (respectively o'_1) be a contact point between the line ℓ_1 (respectively ℓ'_1) and the convex body B . The segment $o_1o'_1$ crosses the segment q_1q_3 and is contained in B by convexity (Figure 21(b)), concluding the proof by contraposition. □

7 Conclusion and Open Problems

Flip graphs of geometric reconfiguration problems are mathematically challenging objects. The flip operation that removes two crossing segments and inserts two non-crossing segments with the same endpoints has applications to local search heuristics and produces directed flip graphs. Such flip graphs have previously been studied in the contexts of TSP tours [37, 43, 46], matchings [10, 12, 21], and trees [10].

In this paper, we have considered several different versions of the problem within the unifying framework of removal and insertion choices. In all versions, we consider a multiset of n segments. When the endpoints are in convex position, we showed that removal choice is enough to get an $O(n \log n)$ bound on the number of flips, implying the same bound for trees (matching the bound in [10]) but without using any specific property of trees. It is an open problem whether

any of these two bounds (the general removal choice bound or the special tree bound⁴) could be improved to $O(n)$ as in the special cases of TSP tours and red-blue matchings. Still in the convex case, we proved an $O(n \log n)$ bound on the number of flips using only insertion choice. We do not know if similar bounds hold in the simple cases of one endpoint inside the convex hull or two points outside the convex hull that cannot be separated by two parallel lines.

The remaining results consider near-convex sets of endpoints. The most general of these results is in the case of no choice, where the endpoints T that are not in convex position may come in any number and may be placed anywhere. In contrast, in our removal strategies, at most two endpoints T that are not in convex position are allowed and the proofs are very technical, considering each placement of T separately. In our insertion strategies, the endpoints T must be outside and separated by two parallel lines. In our removal and insertion strategy the points may be anywhere outside the convex hull but multiple copies of the same segment are not allowed anywhere in the flip graph, which limits its application to matchings or some artificial problems such as all but one vertex having degree 1. Improving and generalizing these bounds remains an open problem.

Our framework with insertion choice is limited to problems where both insertion choices are always available. It is motivating to consider problems where two insertion choices are sometimes available. For example, we may consider the case of maintaining a collection of simple vertex-disjoint graph cycles (without multi-edges). In this case, we have two insertion choices in all but some very specific situations where one of the insertions would create an edge of multiplicity two (that is, a multigraph cycle of length 2). Obtaining a subcubic bound for endpoints in general position with this limited insertion choice is an open problem which may be easier than proving a subcubic bound when insertion choice is never available (with or without removal choice).




Figure 22: A flip sequence where the highlighted segments r_2b_3, r_3b_2 are removed than reinserted and flipped again. It is possible to iterate the use of this same flip $\frac{n}{2}$ times. The first three flips of this figure appear in [12].

Without insertion choice, reducing the gaps between the $O(n^3)$ upper bound and the lower bounds of $\Omega(n^2)$ and $\Omega(n)$ (respectively without and with removal choice) remains the most challenging open problem. The $O(n^{8/3})$ bound on the number of distinct flips is a hopeful step in this direction. Note that there exist flip sequences using the same flip a linear number of times (Figure 22). However, we were not able to find a set of line segments that requires the same pair of segments to be flipped twice in order to be untangled. A proof of this statement would imply a subcubic upper bound with removal choice.

References

- [1] Oswin Aichholzer, Brad Ballinger, Therese Biedl, Mirela Damian, Erik D Demaine, Matias Korman, Anna Lubiw, Jayson Lynch, Josef Tkadlec, and Yushi Uno. Reconfiguration of non-crossing spanning trees. *arXiv preprint*, 2022. URL: <https://arxiv.org/abs/2206.03879>.

⁴A fatal flaw has been found in the proof of the $O(n)$ bound for trees presented in the PhD dissertation [40].

- [2] Oswin Aichholzer, Anna Brötzner, Daniel Perz, and Patrick Schnider. Flips in odd matchings. In *40th European Workshop on Computational Geometry (EuroCG)*, volume 40, pages 447–452, 2024. URL: https://eurocg2024.math.uoi.gr/data/uploads/paper_59.pdf.
- [3] Oswin Aichholzer, Kristin Knorr, Maarten Löffler, Zuzana Masárová, Wolfgang Mulzer, Johannes Obenaus, Rosna Paul, and Birgit Vogtenhuber. Flipping plane spanning paths. In *International Conference and Workshops on Algorithms and Computation (WALCOM)*, 2023. URL: <https://arxiv.org/abs/2202.10831>, doi:10.1007/978-3-031-27051-2_5.
- [4] 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. URL: <https://arxiv.org/abs/1209.0579>, doi:10.1007/s00454-015-9709-7.
- [5] Selim G. Akl, Md Kamrul Islam, and Henk Meijer. On planar path transformation. *Information processing letters*, 104(2):59–64, 2007. doi:10.1016/j.ipl.2007.05.009.
- [6] David L. Applegate, Robert E. Bixby, Vašek Chvátal, and William J. Cook. The traveling salesman problem. In *The Traveling Salesman Problem*. Princeton university press, 2011.
- [7] Sanjeev Arora. Polynomial time approximation schemes for Euclidean TSP and other geometric problems. In *37th Conference on Foundations of Computer Science*, pages 2–11, 1996.
- [8] Sergey Bereg and Hiro Ito. Transforming graphs with the same degree sequence. In *Computational Geometry and Graph Theory*, pages 25–32, 2008. doi:10.1007/978-3-540-89550-3_3.
- [9] Sergey Bereg and Hiro Ito. Transforming graphs with the same graphic sequence. *Journal of Information Processing*, 25:627–633, 2017. doi:10.2197/ipsjjip.25.627.
- [10] Ahmad Biniiaz, Anil Maheshwari, and Michiel Smid. Flip distance to some plane configurations. *Computational Geometry*, 81:12–21, 2019. URL: <https://arxiv.org/abs/1905.00791>.
- [11] Marthe Bonamy, Nicolas Bousquet, Marc Heinrich, Takehiro Ito, Yusuke Kobayashi, Arnaud Mary, Moritz Mühlenhaller, 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. doi:10.4230/LIPICs.MFCS.2019.80.
- [12] Édouard Bonnet and Tillmann Miltzow. Flip distance to a non-crossing perfect matching. *arXiv*, 1601.05989, 2016. URL: <http://arxiv.org/abs/1601.05989>.
- [13] Valentino Boucard, Guilherme D. da Fonseca, and Bastien Rivier. Further connectivity results on plane spanning path reconfiguration. *arXiv*, 2407.00244, 2024. URL: <https://arxiv.org/abs/2407.00244>.
- [14] Nicolas Bousquet and Alice Joffard. Approximating shortest connected graph transformation for trees. In *Theory and Practice of Computer Science*, pages 76–87, 2020. doi:10.1007/978-3-030-38919-2_7.
- [15] Maike Buchin and Bernhard Kilgus. Fréchet distance between two point sets. *Computational Geometry*, 102:101842, 2022. doi:10.1016/j.comgeo.2021.101842.

- [16] Jou-Ming Chang and Ro-Yu Wu. On the diameter of geometric path graphs of points in convex position. *Information processing letters*, 109(8):409–413, 2009. doi:10.1016/j.ipl.2008.12.017.
- [17] Miroslav Chlebík and Janka Chlebíková. Approximation hardness of Travelling Salesman via weighted amplifiers. In *25th International Computing and Combinatorics Conference*, pages 115–127, 2019.
- [18] Nicos Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical report, Carnegie-Mellon Univ Pittsburgh Pa Management Sciences Research Group, 1976.
- [19] Guilherme D. da Fonseca, Yan Gerard, and Bastien Rivier. On the longest flip sequence to untangle segments in the plane. In *International Conference and Workshops on Algorithms and Computation (WALCOM 2023)*, volume 13973 of *Lecture Notes in Computer Science*, pages 102–112, 2023. URL: <https://arxiv.org/abs/2210.12036>, doi:10.1007/978-3-031-27051-2_10.
- [20] Guilherme D. da Fonseca, Yan Gerard, and Bastien Rivier. Short flip sequences to untangle segments in the plane. In *International Conference and Workshops on Algorithms and Computation (WALCOM 2024)*, volume 14549 of *Lecture Notes in Computer Science*, pages 163–178, 2024. URL: <https://arxiv.org/abs/2307.00853>, arXiv:2307.00853.
- [21] Arun Kumar Das, Sandip Das, Guilherme D. da Fonseca, Yan Gerard, and Bastien Rivier. Complexity results on untangling red-blue matchings. *Computational Geometry*, 111:101974, 2023. URL: <https://arxiv.org/abs/2202.11857>, doi:10.1016/j.comgeo.2022.101974.
- [22] Donald Davendra. *Traveling salesman problem: Theory and applications*. BoD–Books on Demand, 2010.
- [23] 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. URL: <https://link.springer.com/content/pdf/10.1007/s00453-013-9801-4.pdf>.
- [24] Paul Erdős, László Lovász, A. Simmons, and Ernst G. Straus. Dissection graphs of planar point sets. In *A survey of combinatorial theory*, pages 139–149. Elsevier, 1973.
- [25] 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. URL: <https://arxiv.org/abs/1205.2842>.
- [26] Branko Grünbaum. Polygons in arrangements generated by n points. *Mathematics Magazine*, 46(3):113–119, 1973. doi:10.1080/0025570X.1973.11976293.
- [27] Gregory Gutin and Abraham P. Punnen. *The traveling salesman problem and its variations*, volume 12. Springer Science & Business Media, 2006.
- [28] 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.
- [29] 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.

- [30] Ferran Hurtado, Marc Noy, and Jorge Urrutia. Flipping edges in triangulations. *Discrete & Computational Geometry*, 22(3):333–346, 1999.
- [31] Alice Joffard. *Graph domination and reconfiguration problems*. PhD thesis, Université Claude Bernard Lyon 1, 2020.
- [32] Linda Kleist, Peter Kramer, and Christian Rieck. On the connectivity of the flip graph of plane spanning paths. In *Graph-Theoretic Concepts in Computer Science (WG 2024)*, 2024. URL: <https://arxiv.org/abs/2407.03912>.
- [33] Charles L. Lawson. Transforming triangulations. *Discrete Mathematics*, 3(4):365–372, 1972.
- [34] Anna Lubiw and Vinayak Pathak. Flip distance between two triangulations of a point set is NP-complete. *Computational Geometry*, 49:17–23, 2015. URL: <https://arxiv.org/abs/1205.2425>.
- [35] Joseph S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems. *SIAM Journal on computing*, 28(4):1298–1309, 1999.
- [36] Naomi Nishimura. Introduction to reconfiguration. *Algorithms*, 11(4), 2018. doi:10.3390/a11040052.
- [37] 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. doi:10.1007/978-3-540-89550-3_17.
- [38] Alexander Pilz. Flip distance between triangulations of a planar point set is APX-hard. *Computational Geometry*, 47(5):589–604, 2014. URL: <https://arxiv.org/abs/1206.3179>.
- [39] Satish B. Rao and Warren D. Smith. Approximating geometrical graphs via “spanners” and “banyans”. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 540–550, 1998.
- [40] Bastien Rivier. *Untangling Segments in the Plane*. PhD dissertation, Université Clermont Auvergne, Clermont-Ferrand, France., November 2023. URL: <https://theses.hal.science/tel-04507187>.
- [41] Jan van den Heuvel. The complexity of change. *Surveys in Combinatorics*, 409:127–160, 2013.
- [42] Jan van Leeuwen and Anneke A. Schoone. Untangling a traveling salesman tour in the plane. Technical Report RUU-CS-80-11, University of Utrecht, 1980. URL: https://dspace.library.uu.nl/bitstream/handle/1874/15912/leeuwen_80_untangling_a_traveling.pdf?sequence=1.
- [43] Jan van Leeuwen and Anneke A. Schoone. Untangling a traveling salesman tour in the plane. In *7th Workshop on Graph-Theoretic Concepts in Computer Science*, 1981.
- [44] Håvard Bakke Bjerkevik, Linda Kleist, Torsten Ueckerdt, and Birgit Vogtenhuber. Flipping non-crossing spanning trees. In *36th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2025. URL: <https://arxiv.org/abs/2410.23809>.
- [45] Todd G. Will. Switching distance between graphs with the same degrees. *SIAM Journal on Discrete Mathematics*, 12(3):298–306, 1999.

- [46] 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.