Université Clermont Auvergne
École Doctoral des Sciences Pour l'Ingénieur
Laboratory of Informatics, Modelling and Optimization of Systems

Untangling Segments in the Plane

PhD Dissertation in Computer Science

Defended by

Bastien Rivier

in November 2023

Jury members	*Affiliation*	*Role*
Guilherme Dias da Fonseca	Aix-Marseille Université	Advisor
Yan Gerard	Université Clermont Auvergne	Advisor
Wolfgang Mulzer	Freie Universität Berlin	Reporter
Carlos Seara	Universidad Politécnica de Catalunya	Reporter
Fatiha Bendali	Université Clermont Auvergne	Examiner
Éric Colin de Verdière	Université Gustave Eiffel	Examiner
Vincent Despré	Université de Lorraine	Examiner
Fabien Feschet	Université Clermont Auvergne	Examiner

Contents

Abstract

How can we reconfigure a given set of n segments with fixed endpoints in the plane into a crossing-free set? We *remove* a pair of crossing segments and *insert* one of the two pairs of non-crossing segments with the same four endpoints in an operation called a *flip*. If we restrict ourselves to sets of segments with a given property, say forming one polygon, the inserted pair must preserve this property. *Untangling* a set of segments amounts to iteratively flip the set of segments until no crossing remains.

Why are flips useful? A flip shortens the total length of the segments. Hence, to efficiently compute approximations of the shortest tour through n given cities, many well known algorithms are supplemented with flips.

How many flips are needed? At most n^3 flips may be performed in a sequence, but the longest sequence known uses only roughly n^2 flips. In general, no strategy for choosing which pairs of crossing segments to remove is known to untangle a set of segments using fewer than n^3 flips. Yet, when the endpoints form a convex polygon, the problem is well understood: the longest sequences use roughly n^2 flips, while the best strategies use roughly n flips.

In this dissertation, we devise strategies to untangle segments for several versions of flips: the flips with nothing to preserve (the segments form a multigraph or a matching), the flips preserving a bipartite matching, the flips preserving a polygon (i.e., a tour), and the flips preserving a tree. We study the performance of each strategy in terms of their number of flips. Our results are organized by the type of choice the strategy uses, as follows. We first study the performance of the strategy choosing nothing, i.e., we improve the bounds on the number of flips in the longest flip sequences (in special cases). We then devise strategies for choosing which pairs of segments to remove to untangle segments using as few flips as possible. We also devise strategies for choosing which pairs of segments to insert, and strategies for choosing both removed and inserted pairs. Many of our results use a parameter measuring how far from a convex polygon the set of endpoints is. Furthermore, we prove the $\mathcal{NP}$-hardness of the shortest flip sequence to untangle a bipartite matching.

Résumé

Comment reconfigurer n segments avec un ensemble fixé d'extrémités pour qu'ils ne se croisent pas ? Nommons *décroisement* l'opération qui consiste à *retirer* une paire de segments qui se croisent pour en *insérer* une autre parmi les deux paires de segments non croisés partageant les mêmes quatre extrémités. Si nous nous restreignons à des ensembles de segments ayant une propriété donnée, à des polygones par exemple, l'insertion de la paire doit préserver cette propriété. *Décroiser totalement* un ensemble de segments consiste à itérer les décroisements jusqu'à ce qu'il ne reste plus de croisement.

À quoi sert de décroiser des segments ? Un décroisement raccourcit la longueur totale des segments. Pour calculer efficacement des approximations du plus court trajet passant par n villes données et revenant à son point de départ, de nombreux algorithmes bien connus font recours à des décroisements pour cette vertu.

Combien de décroisements sont nécessaires pour un décroisement total ? Un décroisement total est toujours atteint en au plus n^3 décroisements, pourtant, la plus longue séquence connue ne comporte qu'environ n^2 décroisements. En général, nous ne savons pas comment choisir stratégiquement quelles paires de segments retirer pour décroiser totalement les ensembles de segments en utilisant moins de n^3 décroisements. Le problème est cependant bien compris lorsque les extrémités des segments forment un polygone convexe : les séquences les plus longues comportent environ n^2 décroisements, tandis que les meilleures stratégies utilisent environ n décroisements.

Dans cette thèse, nous élaborons des stratégies pour décroiser totalement des segments, et ce pour plusieurs versions de décroisements : les décroisements n'ayant rien à préserver (les segments forment par exemple un multigraphe ou un couplage), les décroisements préservant un couplage bipartite, les décroisements préservant un polygone (c'est-à-dire un trajet revenant à son point de départ), et les décroisements préservant un arbre. Nous étudions les performances de chaque stratégie en termes de nombre de décroisements utilisés. Nos résultats sont organisés suivant le type de choix autorisé pour les stratégies. Nous étudions d'abord les performances de la stratégie qui ne choisit rien, c'est-à-dire que nous améliorons les bornes connues sur le nombre de décroisements dans les séquences de décroisements les plus longues (dans certains cas particuliers). Nous élaborons ensuite des stratégies pour choisir les paires de segments à retirer pour décroiser totalement des segments en utilisant le moins de décroisements possible. Nous élaborons également des stratégies pour choisir les paires de segments à insérer, ainsi que des stratégies pour choisir à la fois les paires à retirer et les paires à insérer. Beaucoup de nos résultats utilisent un paramètre mesurant à quel point l'ensemble des extrémités des segments est proche d'un polygone convexe. D'autre part, nous prouvons qu'il est $\mathcal{NP}$-dur de calculer la plus courte séquence de décroisements pour décroiser totalement un couplage bipartite donné.

Chapter 1

Introduction

Motivation. *Imagine*[1] that we have to visit n given cities and come back to the city we first visited (Figure 1.1(a) shows $n = 5$ cities and the roads between). Assuming that there is a flat and straight road between any two cities, what tour should we choose to minimize our total travel distance? This is perhaps the most natural version of the famous Traveling Salesperson Problem (TSP). Computing the exact shortest tour may require an exponential time,[2] which is not practical for several applications (Figure 1.1(b) shows the shortest tour on five cities). To still compute a short tour when the exact shortest tour is not affordable, many powerful heuristics and approximation algorithms have been developed.[3] Many of them might produce self-intersecting tours like the tour in Figure 1.1(c).[4] Yet, a self-intersecting tour is not optimal as it can be reconfigured into a shorter tour by replacing a pair of crossing segments by a non-crossing pair in an operation called a *flip* (a flip reconfigures the tour in Figure 1.1(c) into the tour in Figure 1.1(d) which is shorter as shown in Figure 1.1(e)).[5]

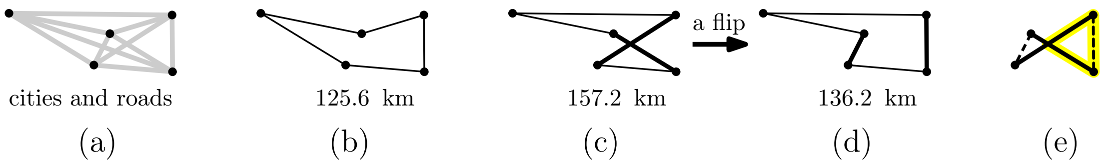

Figure 1.1: (a) Five cities and the roads between. (b) The shortest tour (of 125.6 km). (c) A tour (of 159.3 km) with a crossing. (d) The tour (of 136.2 km) after a flip. (e) In the highlighted triangle, the dashed segment is shorter than the path formed by the two plain segments. The same holds in the other triangle, thus a flip shortens the tour.

[1] The author believes that doing mathematics consists in imagining anything as long as it helps understand something.

[2] The 2-dimension Euclidean TSP is $\mathcal{NP}$-hard [58, 74]. The fastest known algorithm has a worst-case time complexity of $2^{O(\sqrt{n})}$ [35]. A $2^{o(\sqrt{n})}$ time algorithm does not exist unless ETH fails [35, 36].

[3] See for example [9, 34, 51].

[4] A non-exhaustive list of algorithms that might produce self-intersecting tours is: Nearest Neighbor, Greedy, Cheapest Insertion, Nearest Insertion, Farthest Insertion, Random Insertion, Minimum Spanning Tree, Christofides–Serdyukov $\frac{3}{2}$-approximation algorithm [26, 80], all the PTAS [10, 69, 76], and the probabilistic $\frac{3}{2} - \epsilon$-approximation algorithm [62, 78].

[5] This folklore theorem, "Perhaps the oldest theorem concerning the TSP" [29], is true except if the cities are all colinear, as proved in [82]. If general position is assumed, applying the triangle inequality in the two triangles formed by the removed and inserted segments is enough (Figure 1.1(e)).

A reconfiguration problem. This flip operation falls in the range of *reconfiguration problems*, where some typical questions are "Is it possible to reconfigure a chair into a stool while preserving the fact that we can sit on it at every reconfiguration step?" and "What are the minimum number and maximum number of reconfiguration steps if the chair and the stool are chosen to be as different as possible?". We present some related reconfiguration problems in Section 1.4.

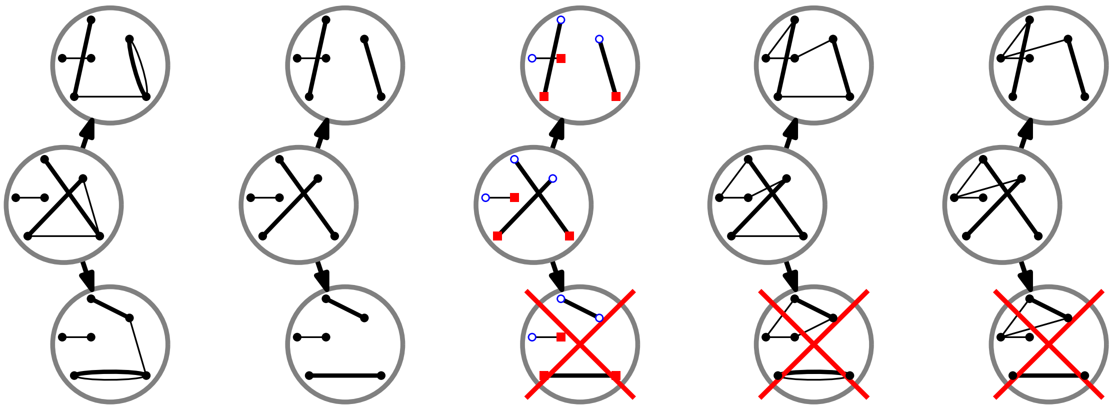

Figure 1.2: The valid insertion choice(s) to preserve (a) nothing special, (b) nothing special (but each point is the end point of exactly *one* segment), (c) segments with *one* red endpoint and *one* blue endpoint, (d) *one* tour (e) *one* tree. Small black discs are points, solid squares are red points and hollow circles are blue points, removed and inserted segments are bold (to be inserted segments are dashed in other figures).

From tours to segments. Flips are not specific to tours: a flip may apply to any set[6] of n segments in the plane. In fact, all the results known for tours are essentially more general. In this generalized setting, performing a flip involves *choosing* which pair of crossing segments to *remove* and which of the two pairs of non-crossing segments with the same four endpoints to *insert*. When flipping a set of segments that has some property, for instance the property of being a tour, we additionally require the resulting set of segments to preserve this property (the same way we required each reconfiguration step of a chair to preserve the property of holding a sitting person). This requirement may impose which of the two possible pairs of non-crossing segments to insert (Figure 1.2(a) shows the two possible insertion choice in an arbitrary (multi)set of segments, Figure 1.2(b) shows the two possible insertion choice in a matching, i.e., a set of segments where each point is the endpoint of exactly one segment, Figure 1.2(c) shows the only insertion choice to preserve a bipartite matching, i.e., a matching where each segment has one red endpoint and one blue endpoint, Figure 1.2(c) shows the only insertion choice to preserve a tour, i.e., a polygon, Figure 1.2(d) shows the only insertion choice to preserve a tree, i.e., a set of roads connecting all the cities but with no cycle). Using the same example of tours, note that one of the two pairs of non-crossing segments with the same four endpoints would transform the tour into two

[6]In fact, multisets of segments must be considered as a flip may insert another copy of an existing segment (Figure 1.2(a)).

tours, which is not wanted. We therefore insert the other pair in order to preserve the property of being *one* tour (Figure 1.2(d)).

Untangling segments. A flip may counter-intuitively increase (or leave unchanged) the total number of crossings; this occurs when some segments of our set cross one of the inserted segments without crossing any of the removed segments (the smallest horizontal segment in Figure 1.2). It is thus unclear whether performing flips as long as there remain crossings eventually terminates. To show that such a process terminates nonetheless with a set of crossing-free segments using at most an exponential number of flips, recall that each flip shortens the total length of the segments (Figure 1.1(e)). It therefore impossible to reach twice the same set of segments in the process. As there is an exponential number of ways n segments can connect the same set of endpoints (Figure 1.3), the claim holds. This *untangling* process is used in practice to supplement heuristics and approximation algorithms for the TSP.[7]

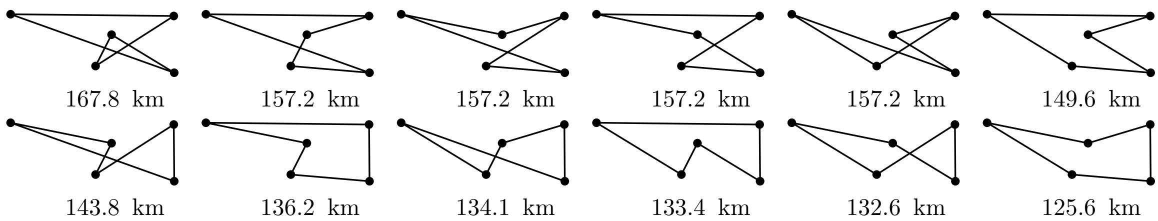

Figure 1.3: All the $\frac{(n-1)!}{2} = 12$ possible tours on the $n = 5$ cities of Figure 1.1(a) sorted by decreasing length. A flip shortens a tour, thus a flip always moves forward in this list.

A historical bound. In 1980, Jan van Leeuwen and Anneke A. Schoone showed that untangling segments uses at most n^3 flips [82],[8] which is much less than an exponential number. This n^3 upper bound is still the state of the art, yet the longest untangling process known only uses roughly n^2 flips[9] which corresponds to the maximum number of crossings in a set of n segments. It is even more surprising that, in general, no strategy guiding the removal choices is known to use fewer flips than n^3 when no set of segments is known to require more than roughly n flips to be untangled, leaving an even wider gap.

Nevertheless, these two gap problems are solved in the very special case where the endpoints of the segments (the cities, in the TSP context) form a *convex* polygon:[10] the longest untangling process uses roughly n^2 flips[11] and the best removal strategy

[7] Untangling is the only known way to get a crossing-free tour shorter than the initial tour in polynomial time [82]. In some situation a crossing-free solution is necessary [22].

[8] This n^3 upper bound, shown for tours, easily generalizes to any multiset of segments.

[9] This $\frac{n(n-1)}{2} = O(n^2)$ lower bound is known for matchings with endpoints forming a convex polygon [19]; the reductions in Theorem 4.1.1 extend it to tours, trees, bipartite matchings, and to general multisets of segments.

[10] In this case, computing the convex-hull of the point set yields the optimal tour in $n \log n$ time [12, 24], so the untangling process seems somewhat useless.

[11] In the case of general multisets of segments, the exact number of flips is $\frac{n(n-1)}{2}$, the number of pairs of segments.

uses roughly n flips. Whether these two bounds also hold for general point sets is an open conjecture[12] from which stems the following problematic.

1.1 Problematic

The problematic of this dissertation is to devise strategies (for various choices) to untangle segments (while preserving various properties) and to establish their performance (mostly in terms of number of flips).

1.2 Contribution

In this section, we present the contributions of this dissertation. These contributions appear in [30, 31, 32, 33] (except Theorem 5.3.1). We start by explaining a notation that we use extensively to present the results. This notation is defined more formally in Chapter 2.

The unknown of this dissertation. Let $\mathbf{d}$ denote the performance of the optimal strategy;[13] it is the unknown of this dissertation. To specify the choice the optimal strategy has, we use symbols in the exponent, as explained in the following.

- In $\mathbf{d}^{\emptyset}$, the optimal strategy has no choice at all, it amounts to choose nothing. The quantity $\mathbf{d}^{\emptyset}$ is thus the maximum number of flips used in an untangling process, and the state of the art is $n^2 \preccurlyeq \mathbf{d}^{\emptyset} \preccurlyeq n^3$ as we have seen in a previous paragraph.[14]
- In $\mathbf{d}^{\mathtt{R}}$, the optimal strategy is a strategy for all the removal choices. The state of the art is $n \preccurlyeq \mathbf{d}^{\mathtt{R}} \preccurlyeq n^3$ as we have seen.
- In $\mathbf{d}^{\mathtt{I}}$, the optimal strategy is a strategy for all the insertion choices.
- In $\mathbf{d}^{\mathtt{RI}}$, the optimal strategy is a strategy for both removal and insertion choices.

We additionally specify the property to be preserve by the flips in the index. For example, in $\mathbf{d}^{\mathtt{R}}_{\mathtt{Cycle}}$, all the sets of segments considered form a cycle (i.e., a tour) and all the flips considered preserve this property. The extensive list of the properties considered in this dissertation is in Chapter 2. We classify these properties into the following three types, and we often combine them.

1. The *point set properties* are the properties about the position of the endpoints that are preserved by any insertion choice at every flip such as `Convex` where the endpoints form a convex polygon (Figure 1.4(a)). We also consider seven different relaxations of this `Convex` property (Figure 1.4(b)-(h), (k)).

[12]The conjecture that the longest untangling process for matchings uses $\Theta(n^2)$ flips appears in [19].
[13]The worst case flip complexity of the problem of untangling segments.
[14]We use Hardy's notation $f(n) \preccurlyeq g(n)$ instead of the Bachmann-Landau notation $f(n) = O(g(n))$ in this chapter to ease the understanding of the non-specialist reader.

2. The *degree properties* are the properties about the segments that are preserved by any insertion choice at every flip.[15] We consider only one such property: `Matching` where each endpoint is matched (via a segment) to exactly one other endpoint (Figure 1.2(b) show flips on a matching).

3. The *insertion properties* are the properties that impose which pair is inserted at *every*[16] flip. We consider three such properties: `Bipartite` where the endpoints are colored red or blue and the segments must match a red point to a blue point (Figure 1.2(c)), `Cycle` where the segments form a tour (Figure 1.2(d)), and `Tree` where the segments form a tree (Figure 1.2(e)). We use `Multigraph` to explicitly say that the insertion choice is not impose by a property (Figure 1.2(a)).

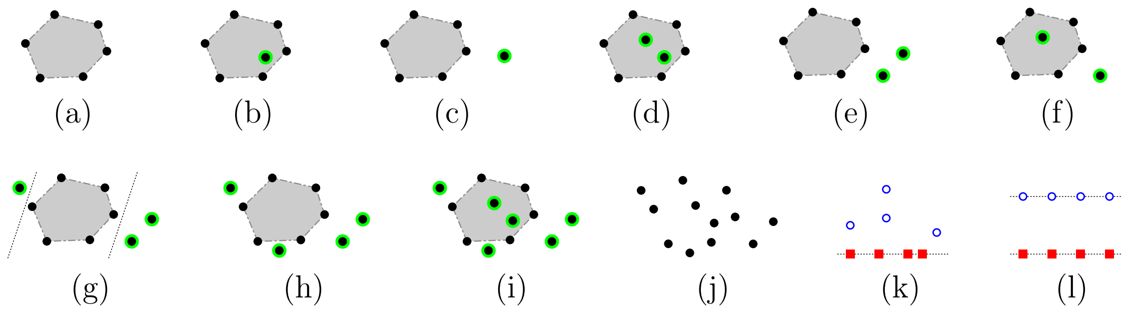

Figure 1.4: The different types of point sets considered in this dissertation. In the nine first figures, the point set P is partitioned into $P = C \cup T$ with C forming a convex polygon (shaded) and the points of T (highlighted) positioned in various ways.
(a) A point set $P = C$ forming a convex polygon (i.e., satisfying the `Convex` property).
(b) & (c) A point set with only one point in T (i.e., satisfying the $|\mathtt{T}| = \mathtt{1}$ property).
(d) A point set with two points in T both lying inside C (i.e., satisfying the `Inin` property).
(e) A point set with two points in T both lying outside C (i.e., satisfying the `Outout` property).
(f) A point set with two points in T, one inside and one outside C (i.e., satisfying the `Inout` property).
(g) A point set where the points of T are separated from C by two parallel lines (i.e., satisfying the `Separated` property).
(h) A point set where the points of T all lie outside C (i.e., satisfying the `Allout` property).
(i) & (j) A general point set (in (i), the point set is viewed as a $P = C \cup T$).
(k) A point set with n red points (solid squares) on a line and n blue points (hollow circles) above the line (i.e., satisfying the `Redonaline` property).
(l) A point set with n red points on a line and n blue points on a parallel line (i.e., satisfying the `Permutation` property).

[15] They are constrains on the degree of each point.

[16] In particular, we do not consider properties such as "admitting a given coloring" where the given coloring has more than two colors.

1.2.1 Upper Bounds

Note that devising a new strategy to untangle segments and measuring its performance upper bounds the corresponding version of $\mathbf{d}$.

Red-on-a-line point sets. The version of the problem where half of the points are red and lie on a line while the other half is blue and lie on a parallel line, and where a segment must match a red point to a blue point (the `Bipartite Matching` version) is fully solved as a set of segments is essentially a drawing of a permutation (Figure 1.4(l)).[17] An intermediate setting where all the red points lie on the same horizontal line and where all the blue points lie above this line (the `Redonaline` version, Figure 1.4(k)) has been studied in [17]. The following two results prolong this study.

In Theorem 4.4.1 [33], we prove that[18]

$$\mathbf{d}^{\emptyset}_{\texttt{Redonaline Matching}}(n) \leq \frac{n(n-1)}{2}\,\frac{n+4}{6} \preccurlyeq n^3,$$

which only betters the state of the art general historic n^3 upper bound [82] by a factor $\frac{1}{6}$. The potential involved in this proof takes advantage of the connections bipartite matchings and permutations, and significantly differs form the potential used to prove the n^3 upper bound.

In Theorem 5.8.1 [33], we provide a simple removal strategy proving that

$$\mathbf{d}^{\mathtt{R}}_{\texttt{Redonaline Matching}}(n) \leq \frac{n(n-1)}{2} \preccurlyeq n^2.$$

This bound is only slightly better than the one in [17], but the removal strategy is simpler and, most importantly, the proof introduces a new approach called *state tracking* which sheds a new light on the untangling problem in general. For instance, state tracking provides a new level of understanding for Theorem 3.2.2 [17] and Theorem 4.2.1 [33].

Convex position point sets. The case where the endpoints of the segments form a convex polygon (Figure 1.4(a)) has a very rigid structure[19] and is thus well understood through a collection bounds, most of which are tight, in the literature. The following three results add to this collection.

In Theorem 5.2.1 [31], we use removal choice and the notion of *crossing depth* of a segment to prove that any n segments with endpoints C forming a convex polygon may be untangled with roughly $n \log n$ flips. Stated as an upper bound, we have[20]

$$\mathbf{d}^{\mathtt{R}}_{\texttt{Convex Multigraph}}(n) \preccurlyeq n \log |C| \preccurlyeq n \log n.$$

[17] Crossings correspond to inversions, flips to inversion swaps, and untangling to sorting.

[18] We purposely factor the bounds by $\frac{n(n-1)}{2} = \binom{n}{2}$, the number of pairs of segments, as this quantity is closely related to the problem.

[19] The geometry of such a convex position is fully encoded in the cyclic order in which the points appear on the convex polygon.

[20] 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 $O(n \log(2 + |C|))$.

This bound was already known for n segments forming a tree [17]; this new proof is both simpler and more general.

In Theorem 5.3.1, we use removal choice and the notion of *crossing depth* again to prove that

$$\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Tree}}(n) \leq 3n - 8 \preccurlyeq n \quad \text{for } n \geq 3,$$

thereby improving the aforementioned $n \log n$ upper bound [17].

In Theorem 6.1.1 [31], we use insertion choice to prove that

$$\mathbf{d}^{\mathtt{I}}_{\mathtt{Convex\ Multigraph}}(n) \preccurlyeq n \log |P| \preccurlyeq n \log n.$$

The proof relies on a multiplicative potential instead of a classic additive potential. This bound improves on the general quadratic upper bound concerning insertion choice [19].

Near convex position point sets. As mentioned, endpoints forming a convex polygon admit a special treatment, but so far no work has been done to take advantage of situations where almost all the endpoints of the segments form a convex polygon (Figure 1.4(b)-(i)). The following eight upper bounds address this situation using the same notations described next. The set of endpoints is denoted P and is partitioned into two subsets: the subset C forming a convex polygon and the subset T of all the remaining endpoints. We introduce the parameter t defined as the number of segments which have at least an endpoint in T, counting twice a segment having its two endpoints in T. This parameter t is a measure of how "far" the point set P is from being in convex position with C. The proofs often involve a preprocessing of the segments whose endpoints are both in C invoking one of the corresponding theorems for convex position. The proofs also involve *splitting* the segments to untangle them separately (see Lemma 2.6.1 [19, 31]).

In Theorem 4.3.1 [30], we prove that

$$\mathbf{d}^{\emptyset}_{\mathtt{Multigraph}}(n, t) \preccurlyeq tn^2.$$

This bound consists in a continuous transition between the quadratic bound from Theorem 3.2.2 [17] when $t = 0$ and the cubic bound from Theorem 3.1.3 [82] when $t = n$. The proof indeed combines the potentials from both proof.

In Theorem 5.4.2 [31], we provide a removal strategy for the case where there is only one point in T to prove that

$$\mathbf{d}^{\mathtt{R}}_{|\mathtt{T}|=1\ \mathtt{Multigraph}}(n, t) \preccurlyeq n \log |C| + tn \preccurlyeq n \log n + tn.$$

In Theorem 5.5.2 [31], we provide a removal strategy for the case where there are two points in T, one inside and one outside the convex polygon formed by C, to prove that

$$\mathbf{d}^{\mathtt{R}}_{\mathtt{Inout\ Multigraph}}(n, t) \preccurlyeq t^2 n + n \log n.$$

In Theorem 5.6.1 [31], we provide a removal strategy for the case where there are two points in T, both lying inside the convex polygon formed by C, to prove that

$$\mathbf{d}^{\mathtt{R}}_{\mathtt{Inin\ Multigraph}}(n, t) \preccurlyeq tn + n \log n.$$

In Theorem 5.7.1 [31], we provide a removal strategy for the case where there are two points in T, both lying outside the convex polygon formed by C, to prove that

$$\mathbf{d}^{\mathtt{R}}_{\mathtt{Outout\ Multigraph}}(n,t) \preccurlyeq 2^t n \log n.$$

In Theorem 6.2.1 [31], we provide an insertion strategy for the case where the points of C lie in between two parallel lines while the points of T lie outside this strip is only one point in T to prove that

$$\mathbf{d}^{\mathtt{I}}_{\mathtt{Separated\ Multigraph}}(n,t) \preccurlyeq t\,|P| \log |C| + n \log |C| \preccurlyeq tn \log n.$$

In Theorem 7.1.1 [31], we provide a strategy for both removal and insertion choices for the same case as Theorem 6.2.1 [31] (described in the previous paragraph) to prove that

$$\mathbf{d}^{\mathtt{RI}}_{\mathtt{Separated\ Multigraph}}(n,t) \preccurlyeq n + t\,|P| \preccurlyeq tn.$$

In Theorem 7.2.3 [31], we provide a strategy for both removal and insertion choices for the case where all the points of T lie outside the convex polygon formed by C and where no endpoint is shared by more than one segment, to prove that

$$\mathbf{d}^{\mathtt{RI}}_{\mathtt{Allout\ Matching}}(n,t) \preccurlyeq t^3 n.$$

Counting flips without multiplicity. In all the previous upper bounds, using many times the same flip had to be counted accordingly in the performance. In Theorem 4.5.1 [30], we prove the sub-cubic $n^{8/3}$ upper bound on the number of *distinct* flips, that is, the number of flips where a flip used many times only count for one. This upper bound may become useful if combined with a strategy avoiding to use twice the same flip.

1.2.2 Lower Bounds

Lower bounds are very important as they are the only hint of how far our strategies may be from being optimal. In this dissertation, we show the following two lower bounds.

In Theorem 4.2.1 [33], we show that

$$n^2 \preccurlyeq \frac{3}{2}\,\frac{n(n-1)}{2} - \frac{n}{4} \leq \mathbf{d}^{\emptyset}_{\mathtt{Redonaline\ Matching}}(n).$$

Previously to this result, the longest untangling process known used $\frac{n(n-1)}{2}$ flips and had the endpoints of the segments forming a convex polygon. This new lower bound, even though roughly similar, disproves the intuition that the longest untangling process occurs when the endpoints of the segments form a convex polygon.

In Theorem 5.1.1 [33], we show that

$$n \preccurlyeq \frac{3}{2}n - 2 \leq \mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Bipartite\ Matching}}(n).$$

It is the first lower bound on $\mathbf{d}^{\mathtt{R}}$ that is greater than n.

1.2.3 Reductions

All these bounds may seem very specialized to a specific version of $\mathbf{d}$ at first sight, but most of them carry over to a number of other version of $\mathbf{d}$ using the following four reduction results. Lemma 2.3.1 essentially states that the more choice we have, the shorter the untangling process. Lemma 2.3.2 essentially states that the stronger the property to be preserved, the shorter the untangling process. Lemma 2.3.3 roughly states that preserving a weaker property in exchange of more choice may lead to shorter untangling processes. Theorem 4.1.1 [30] implies that essentially all the versions of $\mathbf{d}^{\emptyset}$ are roughly the same. We highlight our contributions and summarize all the upper and lower bounds on all the versions of $\mathbf{d}$, including all the bounds derived thanks to reductions (the original theorems are marked in bold), in Table 1.1 and Table 1.2 for the `Multigraph` versions, in Table 1.3 and Table 1.4 for the `Matching` versions, in Table 1.5 for the `Bipartite Matching` versions, in Table 1.6 for the `Cycle` versions, and in Table 1.7 for the `Tree` versions. All these tables are in Section 1.6 at the end of this chapter.

1.2.4 Intractability

All the contributions presented so far are related to $\mathbf{d}$, i.e. to counting flips. In contrast, this last contribution is about counting time. In the version of the problem where half of the points are red and the other half is blue, and where a segment must match a red point to a blue point, we show in Theorem 8.0.1 that it is not possible (unless $\mathcal{P} = \mathcal{NP}$) to compute the shortest possible untangle process of a given set of segments within a reasonable (i.e., polynomial) time.

1.3 Organization

This dissertation is organized into nine chapters.

Chapter 1 is the present introduction, where we first introduce the problematic, mention the contributions, describe the organization, then we present related reconfiguration problems.

Chapter 2 states the assumptions and provides definitions and lemmas used intensively throughout the dissertation.

Chapter 3 presents the literature results falling in the problematic of this dissertation, providing proofs and generalisations for most of them. In contrast with the general organization by choice, the sections of Chapter 3 are organized by point set property, highlighting the fact that choice did not benefit from a systematic treatment until now.

The rest of the chapters presents the contributions of this dissertation.

Chapter 4 is dedicated to untangle sequences with no choice. We first prove that all the versions of $\mathbf{d}^{\emptyset}$ have the same asymptotic behaviour. Then we prove one lower bound and two upper bounds on various versions of $\mathbf{d}^{\emptyset}$. Finally, we prove a sub-cubic upper bound on the number of *distinct* flips in an untangle sequence.

Chapter 5 is dedicated to untangle sequences with removal choice. We prove one lower bound and seven upper bounds on various versions of $\mathbf{d}^{\mathtt{R}}$.

Chapter 6 is dedicated to untangle sequences with insertion choice. We prove two upper bounds on two different versions of $\mathbf{d}^{\mathtt{I}}$.

Chapter 7 is dedicated to untangle sequences with both removal and insertion choices. We prove two upper bounds on two different versions of $\mathbf{d}^{\mathtt{RI}}$.

Chapter 8 is dedicated to proving that any constant factor approximation of the length of the shortest untangle sequence in the `Bipartite Matching` version is $\mathcal{NP}$-hard.

Chapter 9 contains some elements of discussion about the open problems raised by the problematic of the dissertation. In particular, we give some insight on the difficulties to extend known results, we provide counter-examples to some ideas, and we discuss about what we believe to be interesting ideas to tackle some of the open problems.

1.4 Related Reconfiguration Problems

Untangling segments is one of a wide range of problems which share a common setting: the reconfiguration problems. A reconfiguration problem always comes with a graph, called the *reconfiguration graph*, where the vertices are combinatorial or geometrical objects called the *configurations* (in our case, the multisets of segments satisfying a given property on a given point set), and where there is a directed edge from a configuration to another if there exists a *reconfiguration* (in our case, a flip) transforming the first into the second (Figure 1.5). Some typical questions are to find whether this reconfiguration graph is connected, to find its diameter and the length of the longest path (for an overview on reconfiguration see for example [20, 60, 81]). Reconfiguration may be used to progressively improve the quality of a solution to a given problem (like in this dissertation, to shorten the length of a tour) or to enumerate efficiently all the solutions of a problem (see for example [54] for an enumeration of crossing-free tours with special visibility properties). In this section, we present some reconfiguration problems related to ours.

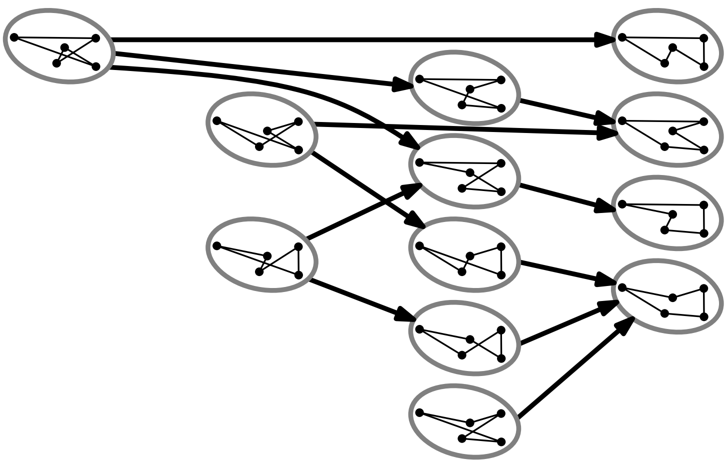

Figure 1.5: The reconfiguration graph of the point set from Figure 1.1(a) for flips in the `Cycle` version. The tours displayed in the same column have the same number of crossings, and are sorted by decreasing length.

Opt reconfiguration. One could wonder about also allowing a non-crossing pair of segments to be replaced as long as the total length of the segments is shortened. This relaxed operation, called a 2-*opt change*, is used in a class of approximation algorithms for the TSP. Note that both flips and 2-opt changes are special cases of 2-changes. A 2-*change* has the same definition as a flip except that the condition that the removed segments cross is dropped (no geometry is involved anymore).

The apparent similarity with flips is misleading. A striking contrast is that there exist sequences with an exponential number of 2-opt changes [42].

Moreover, 2-opt changes depend only on the order of the pairs of segments with endpoints in P sorted by increasing total length (induced by the metric on P) whereas flips depend only on the set of crossing pairs of segments with endpoint in P (called the *crossing type* or the x-*type* of P in the literature [73], the crossing type of P is induced by the order type of P). In more concrete terms, a perturbation of the coordinates may turn a 2-opt change into a 2-change which is not a 2-opt change while preserving the set of crossing pairs of segments with endpoint in P (and the order type of P). We discuss about crossing types in Section 1.5.

As it is a fortiori the case with flips, it is not enough to find an optimal tour ([61], page 18).

The length of a 2-optimal tour in a metric TSP instance with n cities is at most $\sqrt{n/2}$ times the length of a shortest tour and this bound is tight. The approximation ratio of the 2-Opt Heuristic for the metric TSP is $\sqrt{n/2}$ [56].

Triangulation reconfiguration. Reconfiguration problems in the context of triangulations are widely studied [71]. A reconfiguration consists of removing one edge and adding another one while preserving a triangulation. It is known that $\Theta(n^2)$ reconfigurations are sufficient and sometimes necessary to obtain a Delaunay triangulation [57, 64]. Determining the reconfiguration distance between two triangulations of a point set [68, 75] and between two triangulations of a simple polygon [6] are both $\mathcal{NP}$-hard.

Planar tree reconfiguration. The planar tree reconfigurations are similarly to triangulation reconfigurations. A reconfiguration consists of removing one edge and adding another one while preserving a tree [3].

Planar path, a special class of planar trees, have also been studied. It is possible to reconfigure any two non-crossing paths if the points are in convex position [8, 23] or if there is *one* point inside the convex hull [4].

Multigraph reconfiguration. It is possible to relax the flip definition to all operations that replace two segments by two others with the same four endpoints, whether they cross or not, and generalize the configurations to multigraphs with the same degree sequence [52, 53, 60]. In this context, finding the shortest path from a given configuration to another in the reconfiguration graph is $\mathcal{NP}$-hard, yet 1.5-approximable [14, 15, 45, 84]. If we additionally require the configurations to be connected graphs, the same problem is $\mathcal{NP}$-hard and 2.5-approximable [21].

Considering perfect matchings of an arbitrary graph (instead of the complete graph), a flip amounts to exchanging the edges in an alternating cycle of length four. It is then PSPACE-complete to decide whether there exists a path from a configuration to another [18]. There is, actually, a wide variety of reconfiguration contexts derived from $\mathcal{NP}$-complete problems where this same accessibility problem is PSPACE-complete [59].

1.5 Contextual Problems

In this section, we gather some pieces of answer to some relevant questions about untangling segments.

Searching for crossings. Before performing a flip on a multiset S of n segments, we have to search for crossing pairs of segments. The time taken by this search is not taken into account in our flip count; it is nonetheless important to mention the time (and space) complexity of this search.

To find all the crossings of a multiset of segments, we may use Ivan J. Balaban's algorithm [11] asymptotically optimal both in time, $O(k+n\log n)$ where k is the number of intersecting pairs of segments, and space, $O(n)$.[21] In practice, the Bentley–Ottmann sweep line algorithm [13] is often preferred because it is easier to implement and its space complexity is $O(n)$, even though its time complexity $O((k+n)\log n)$ is not asymptotically optimal.[22]

If we are not interested in choosing a crossing pair, but only in finding one, then the Shamos–Hoey algorithm [77] running in $O(n\log n)$ time is enough.[23]

However, if $O(nm) = O(n^2)$ space is not too much to ask (m being the maximum number of crossings appearing on the same segment of S during the untangling process), then we may maintain for each segment $s \in S$ a list of all the segments crossing s. This list is initialized, say, using Ivan J. Balaban's algorithm. Then, at each flip, we remove two segments together with their list of crossings, and we insert two segments and build their list of crossings in $O(n)$. This process only adds a $O(n)$ time overhead at each flip and a preprocessing time of $O(k+n\log n) = O(n^2)$ (where k is the number of crossings in the initial multiset S).

This considerations about the time (and space) overhead induced by the search for crossings does not include the time (and space) overhead induced by the computation of removal and insertion choices which depends on the strategy used. All the insertion strategies presented in this dissertation run in constant time and space per flip. All the removal strategies and all the strategies for both removal and insertion presented in this dissertation run in at most quadratic time (to search for the right crossing pair to remove among the crossings we have computed) and constant space per flip. This rough upper bound may be improved by analyzing each strategy.

[21]Other algorithms are asymptotically optimal in time, but not in space: the randomized ones are [28, 70], and the deterministic one is [25].

[22]Two other radomized algorithms, both based on the Bentley–Ottmann algorithm, additionaly require the segments to form a connected graph: one runs in $O(k+n\log^* n)$ expected time [27], the other runs in $O(n+k\log^{(i)} n)$ expected time for any constant i [43]. In our context, this requirement is fulfilled for cycles and trees.

[23]This algorithm is in fact the precursor of the Bentley–Ottmann algorithm.

Computing crossing-free segments without flips. We mentioned that untangling a tour remains the only known way to compute in polynomial time a crossing-free tour of shorter length. However, it is not the case for matchings, bipartite matchings, and trees where faster algorithms are available. Next, we mention several of these computational geometry algorithms (for an introduction to computational geometry, see for instance [16]).

A simple sweep line algorithm is enough to compute a crossing-free matching on $2n$ points[24] in $O(n \log n)$ time: sort the points by increasing x coordinates and connect every other point to the next one. Notice that the minimal-weight perfect matching computed in $O(n^3)$ time [47] by Edmonds' blossom algorithm [40, 41] is also crossing free.[25]

Given a bipartite point set of $2n$ points,[24] Edmonds' blossom algorithm (in the bipartite complete graph instead of the complete graph) also yields a crossing-free bipartite matching in $O(n^3)$ time. Yet, such a crossing-free bipartite matching can be computed in optimal $O(n \log n)$ time and optimal $O(n)$ space [55]. The algorithm is based on semi-dynamic convex hull data structures and does not use flips. Note that recursively computing a ham-sandwich cut of a bipartite point set of $2n$ points in optimal $O(n)$ time [66, 67] also yields an optimal $O(n \log n)$ time algorithm to compute a crossing-free bipartite matching.[26]

The 2-dimension Euclidean minimum spanning tree of $n + 1$ points[24] can also be computed in optimal $O(n \log n)$ time: compute the Delaunay triangulation in $O(n \log n)$ time [38, 50, 65], then run a graph minimum spanning tree algorithm, say Kruskal's algorithm in $O(m \log n)$ time (where m is the number of edges) [63], on this Delaunay triangulation graph with its $O(n)$ edges weighted with their length. However, we do not have control on the degree of each point of the minimum spanning tree. In fact, computing the spanning tree of minimum total distance with the degree of each point imposed is $\mathcal{NP}$-hard (by reduction of the Hamiltonian path problem). To compute a crossing-free spanning tree on a point set where the degree of each point is imposed, the untangling process may be a good option.

From point sets to crossing types. We have presented the set P of endpoints of the segments as a part of the input of the untangling problem, but the actual coordinates of the points of P are not required if the set of all the crossing pairs of segments between any two points of P is provided. More precisely, we say that two sets of labeled points P_1 and P_2 have the same *labeled crossing type* if their set of crossing pairs of segments both corresponds to the same set of unordered pairs of unordered pairs of labels[27]. We say that two sets of points P_1 and P_2 have the same *crossing type*[28] if there exist labels for each of P_1 and P_2 such that P_1 and P_2 have the same labeled crossing type.

[24] The variable n always refers to the number of segments.

[25] This algorithm is used as a subroutine in Christofides–Serdyukov approximation algorithm for the metric TSP.

[26] The problem of computing crossing-free matchings has also been considered in higher dimensions [7]. The problem of computing minimal-weight perfect matching has been considered on arbitrary bipartite graphs instead of the complete bipartite graph [1].

[27] This repetition is not a typographic error, a segment being represented by an unordered pair of labels.

[28] Crossing types are also called x-types in the literature.

The untangling problem depends only on the crossing type of the point set. Therefore, properties such as having endpoints in convex position should be thought of as having endpoints with the same crossing type as a point set in convex position.

The crossing type of a point set P may be derived from the *order type* of P, which is defined similarly to the crossing type, replacing "the set of all the crossing pairs of segments between any two points of P" by "the set of all the orientations of any three points of P" (three labeled points have three possible orientations: clockwise, counterclockwise, and colinear but this last is excluded by our general position assumption). Crossing types and order types benefit a wide range of applications but raise difficult problems such as counting, enumerating, sampling, finding (integer) coordinate representations, characterizing. They are actively studied (for example see [5] for a reconstruction of all the order types of a given list of cyclic orders, see [73] for an algorithm reconstructing the crossing type of a point set from its planar spanning trees organised in a reconfiguration graph). Order types have been enumerated up to 11 points [2]. However, it does not seem practical to test conjectures from our untangling problem using these first enumerated order types.

1.6 Summary Tables

This section contains tables summarizing all the bounds on all the versions of $\mathbf{d}$. Each cell in the tables contains an asymptotic lower or upper bound on $\mathbf{d}_{\Pi_1\ \Pi}^{\mathit{Choices}}(n)$ for *Choices* $\in \{\emptyset, \mathtt{R}, \mathtt{I}, \mathtt{RI}\}$ specified in the column, for Π_1 specified in the line, and for Π specified in the title of the table, with a reference to the corresponding theorem just bellow. Sometimes an asymptotic bound follows from more than one theorems, in which case they are all mentioned.

We use normal font for the bounds which are only a corollary of the theorems specified via some reductions (Lemma 2.3.1, Lemma 2.3.2, Lemma 2.3.3, and Theorem 4.1.1 [30]). We use bold font for the theorem bounds themselves. When an asymptotic bound is both a theorem and a corollary of another theorem, both theorems are mentioned but only the former is bold. A cell is highlighted when a contribution is involved (i.e., when the theorem proving the cell bound with the best constants is a contribution).

The reductions often transfer the upper bound of a given cell to the upper bound cells on its left, to the upper bound cells bellow, and to the corresponding upper bound cell in a table bellow. Conversely, the reductions often transfer the lower bound of a given cell to the lower bound cells on its right, to the lower bound cells above, and to the corresponding lower bound cell in a table above.

We recall the notation for the different types of choice.

$\emptyset$: No choice, which means the longest untangle sequence.

$\mathtt{R}$: Removal choice.

$\mathtt{I}$: Insertion choice.

$\mathtt{RI}$: Both removal and insertion choices.

We also recall the notation for the different types of point sets.

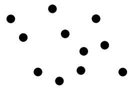

`General`: The endpoints may be anywhere.

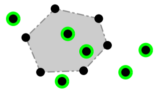

$\mathtt{C} \cup \mathtt{T}$: The endpoints are partitioned into a set C in convex position and a general set T.

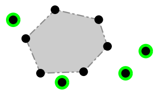

`Allout`: The endpoints are partitioned into a set C in convex position and a set T outside the convex hull of C.

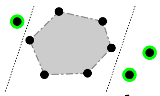

`Separated`: The endpoints are partitioned into a set C in convex position and a set T outside the convex hull of C such that T and C are separated by two parallel lines.

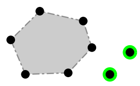

`Outout`: The endpoints are partitioned into a set C in convex position and a set T containing two points outside the convex hull of C.

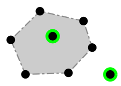

`Inout`: The endpoints are partitioned into a set C in convex position and a set T with one point inside and one point outside the convex hull of C.

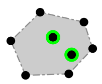

`Inin`: The endpoints are partitioned into a set C in convex position and a set T containing two points inside the convex hull of C.

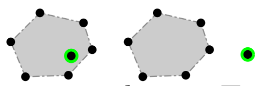

$|\mathtt{T}| = \mathtt{1}$: The endpoints are partitioned into a set C in convex position and a set T containing one point.

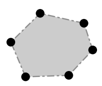

`Convex`: The endpoints are in convex position.

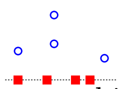

`Redonaline`: The red endpoints are on a line and the blue points are above this line.

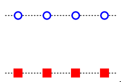

`Permutation`: The endpoints are on two parallel lines, one containing red points and the other blue points.

1.6.1 Multigraph

Multigraph Version with Both Choices and Insertion Choice

Table 1.1: Asymptotic lower and upper bounds on $\mathbf{d}^{Choices}_{\Pi\ \texttt{Multigraph}}$ for the property Π of the point set $P = C \cup T$ specified in each line and for the *Choices* is specified in each column.

	RI		I	
	Lower	Upper	Lower	Upper
`General`	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]	n 3.2.12 [19]	$\boldsymbol{n^2}$ **and** $\boldsymbol{n\sigma(P)}$ **3.1.4 [19] and 3.1.5 [17]**
`C` $\cup$ `T`	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]
`Allout`	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]
`Separated`	n 3.2.12 [19]	$\boldsymbol{tn}$ **7.1.1 [31]**	n 3.2.12 [19]	$\boldsymbol{tn \log n}$ **6.2.1 [31]**
`Outout`	n 3.2.12 [19]	tn 7.1.1 [31]	n 3.2.12 [19]	$tn \log n$ 6.2.1 [31]
`Inout`	n 3.2.12 [19]	$t^2 n$ 5.5.2 [31] using 3.2.13 [17]	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]
`Inin`	n 3.2.12 [19]	tn 5.6.1 [31] using 3.2.13 [17]	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]
$\lvert$`T`$\rvert = 1$	n 3.2.12 [19]	tn 5.4.2 [31] using 3.2.13 [17]	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]
`Convex`	n 3.2.12 [19]	$\boldsymbol{n}$ **3.2.13 [17]**	n 3.2.12 [19]	$\boldsymbol{n \log n}$ **6.1.1 [31]**

Multigraph Version with Removal Choice and No Choice

Table 1.2: Asymptotic lower and upper bounds on $\mathbf{d}^{Choices}_{\Pi\ \texttt{Multigraph}}$ for the property Π of the point set $P = C \cup T$ specified in each line and for the *Choices* is specified in each column.

	R		$\emptyset$	
	Lower	Upper	Lower	Upper
General	n 3.2.12 [19], 5.1.1 [33]	n^3 3.1.3 [19, 82]	n^2 3.2.1 [19], 4.2.1 [33]	$\boldsymbol{n^3}$ **3.1.3 [19, 82]**
C ∪ T	n 3.2.12 [19], 5.1.1 [33]	tn^2 4.3.1 [30]	n^2 3.2.1 [19], 4.2.1 [33]	$\boldsymbol{tn^2}$ **4.3.1 [30]**
Allout	n 3.2.12 [19], 5.1.1 [33]	tn^2 4.3.1 [30]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Separated	n 3.2.12 [19], 5.1.1 [33]	tn^2 4.3.1 [30]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Outout	n 3.2.12 [19], 5.1.1 [33]	$\boldsymbol{2^t n \log n}$ **5.7.1 [31] using 5.2.1 [31]**	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Inout	n 3.2.12 [19], 5.1.1 [33]	$\boldsymbol{t^2 n + n \log n}$ **5.5.2 [31] using 5.2.1 [31]**	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Inin	n 3.2.12 [19], 5.1.1 [33]	$\boldsymbol{tn + n \log n}$ **5.6.1 [31] using 5.2.1 [31]**	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
$\lvert T \rvert = 1$	n 3.2.12 [19], 5.1.1 [33]	$\boldsymbol{tn + n \log n}$ **5.4.2 [31] using 5.2.1 [31]**	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Convex	n 3.2.12 [19], 5.1.1 [33]	$\boldsymbol{n \log n}$ **5.2.1 [31]**	n^2 3.2.1 [19]	$\boldsymbol{n^2}$ **3.2.2 [17]**

1.6.2 Matching

Matching Version with Both Choices and Insertion Choice

Table 1.3: Asymptotic lower and upper bounds on $\mathbf{d}^{\text{Choices}}_{\Pi\ \texttt{Matching}}$ for the property Π of the point set $P = C \cup T$ specified in each line and for the *Choices* is specified in each column.

	RI		I	
	Lower	Upper	Lower	Upper
General	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]
$\mathtt{C} \cup \mathtt{T}$	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]
Allout	n 3.2.12 [19]	$\boldsymbol{t^3 n}$ **7.1.1 [31]**	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]
Separated	n 3.2.12 [19]	tn 7.1.1 [31]	n 3.2.12 [19]	$tn \log n$ 6.2.1 [31]
Outout	n 3.2.12 [19]	n 7.1.1 [31]	n 3.2.12 [19]	$n \log n$ 6.2.1 [31]
Inout	n 3.2.12 [19]	n 5.5.2 [31] using 3.2.13 [17]	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]
Inin	n 3.2.12 [19]	n 5.6.1 [31] using 3.2.13 [17]	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]
$\lvert\mathtt{T}\rvert = 1$	n 3.2.12 [19]	n 5.4.2 [31] using 3.2.13 [17]	n 3.2.12 [19]	n^2 and $n\sigma(P)$ 3.1.4 [19] and 3.1.5 [17]
Convex	n 3.2.12 [19]	n 3.2.13 [17]	n 3.2.12 [19]	$n \log n$ 6.1.1 [31]

Matching Version with Removal Choices and No Choice

Table 1.4: Asymptotic lower and upper bounds on $\mathbf{d}^{\text{Choices}}_{\Pi\ \mathtt{Matching}}$ for the property Π of the point set $P = C \cup T$ specified in each line and for the *Choices* is specified in each column.

	R		∅	
	Lower	Upper	Lower	Upper
General	n 3.2.12 [19], 5.1.1 [33]	n^3 3.1.3 [19, 82]	n^2 3.2.1 [19], 4.2.1 [33]	n^3 3.1.3 [19, 82]
C ∪ T	n 3.2.12 [19], 5.1.1 [33]	tn^2 4.3.1 [30]	n^2 3.2.1 [19], 4.2.1 [33]	tn^2 4.3.1 [30]
Allout	n 3.2.12 [19], 5.1.1 [33]	tn^2 4.3.1 [30]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Separated	n 3.2.12 [19], 5.1.1 [33]	tn^2 4.3.1 [30]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Outout	n 3.2.12 [19], 5.1.1 [33]	$n \log n$ 5.7.1 [31] using 5.2.1 [31]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Inout	n 3.2.12 [19], 5.1.1 [33]	$n \log n$ 5.5.2 [31] using 5.2.1 [31]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Inin	n 3.2.12 [19], 5.1.1 [33]	$n \log n$ 5.6.1 [31] using 5.2.1 [31]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
|T| = 1	n 3.2.12 [19], 5.1.1 [33]	$n \log n$ 5.4.2 [31] using 5.2.1 [31]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Convex	n 3.2.12 [19], 5.1.1 [33]	$n \log n$ 5.2.1 [31]	n^2 3.2.1 [19]	n^2 3.2.2 [17]

1.6.3 Bipartite Matching

Table 1.5: Asymptotic lower and upper bounds on $\mathbf{d}^{\text{Choices}}_{\Pi\ \texttt{Bipartite Matching}}$ for the property Π of the point set $P = C \cup T$ specified in each line and for the *Choices* is specified in each column.

	R		∅	
	Lower	Upper	Lower	Upper
General	n 3.2.12 [19], 5.1.1 [33]	n^3 3.1.3 [19, 82]	n^2 3.2.1 [19], 4.2.1 [33]	n^3 3.1.3 [19, 82]
Redonaline	n 3.2.12 [19], 5.1.1 [33]	$\boldsymbol{n^2}$ 3.3.1 [17], **5.8.1 [33]**	$\boldsymbol{n^2}$ 3.2.1 [19], **4.2.1 [33]**	$\boldsymbol{n^3}$ 3.1.3 [19, 82], **4.4.1 [33]**
C ∪ T	n 3.2.12 [19], 5.1.1 [33]	tn^2 4.3.1 [30]	n^2 3.2.1 [19], 4.2.1 [33]	tn^2 4.3.1 [30]
Allout	n 3.2.12 [19], 5.1.1 [33]	tn^2 4.3.1 [30]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Separated	n 3.2.12 [19], 5.1.1 [33]	tn^2 4.3.1 [30]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Outout	n 3.2.12 [19], 5.1.1 [33]	n 5.7.1 [31] using 3.2.10 [17]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Inout	n 3.2.12 [19], 5.1.1 [33]	n 5.5.2 [31] using 3.2.10 [17]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Inin	n 3.2.12 [19], 5.1.1 [33]	n 5.6.1 [31] using 3.2.10 [17]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
$\lvert T\rvert = 1$	n 3.2.12 [19], 5.1.1 [33]	n 5.4.2 [31] using 3.2.10 [17]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Convex	$\boldsymbol{n}$ 3.2.12 [19], **5.1.1 [33]**	$\boldsymbol{n}$ **3.2.10 [17]**	n^2 3.2.1 [19]	n^2 3.2.2 [17]
Permutation	$\boldsymbol{n}$ **3.2.12 [19]**	n 3.2.10 [17]	$\boldsymbol{n^2}$ **3.2.1 [19]**	n^2 3.2.2 [17]

1.6.4 Cycle

Table 1.6: Asymptotic lower and upper bounds on $\mathbf{d}_{\Pi\ \mathtt{Cycle}}^{\text{Choices}}$ for the property Π of the point set $P = C \cup T$ specified in each line and for the *Choices* is specified in each column.

	R		∅	
	Lower	Upper	Lower	Upper
General	n 3.2.12 [19]	n^3 3.1.3 [19, 82]	n^2 3.2.1 [19], 4.2.1 [33]	n^3 3.1.3 [82]
C ∪ T	n 3.2.12 [19]	tn^2 4.3.1 [30]	n^2 3.2.1 [19], 4.2.1 [33]	tn^2 4.3.1 [30]
Allout	n 3.2.12 [19]	tn^2 4.3.1 [30]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Separated	n 3.2.12 [19]	tn^2 4.3.1 [30]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Outout	n 3.2.12 [19]	n 5.7.1 [31] using 3.2.7 [72], 3.2.9 [85]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Inout	n 3.2.12 [19]	n 5.5.2 [31] using 3.2.7 [72], 3.2.9 [85]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Inin	n 3.2.12 [19]	n 5.6.1 [31] using 3.2.7 [72], 3.2.9 [85]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
$\vert T\vert = 1$	n 3.2.12 [19]	n 5.4.2 [31] using 3.2.7 [72], 3.2.9 [85]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Convex	$\boldsymbol{n}$ 3.2.12 [19]	$\boldsymbol{n}$ 3.2.7 [72], **3.2.9 [85]**	n^2 3.2.1 [19]	n^2 3.2.2 [17]

1.6.5 Tree

Table 1.7: Asymptotic lower and upper bounds on $\mathbf{d}_{\Pi\ \mathtt{Tree}}^{\mathrm{Choices}}$ for the property Π of the point set $P = C \cup T$ specified in each line and for the *Choices* is specified in each column.

	R		∅	
	Lower	Upper	Lower	Upper
General	n 3.2.12 [19]	n^3 3.1.3 [19, 82]	n^2 3.2.1 [19], 4.2.1 [33]	n^3 3.1.3 [82]
C ∪ T	n 3.2.12 [19]	tn^2 4.3.1 [30]	n^2 3.2.1 [19], 4.2.1 [33]	tn^2 4.3.1 [30]
Allout	n 3.2.12 [19]	tn^2 4.3.1 [30]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Separated	n 3.2.12 [19]	tn^2 4.3.1 [30]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Outout	n 3.2.12 [19]	$2^t n$ 5.7.1 [31] using 5.3.1 [31]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Inin	n 3.2.12 [19]	tn 5.6.1 [31] using 5.3.1 [31]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Inout	n 3.2.12 [19]	$t^2 n$ 5.5.2 [31] using 5.3.1 [31]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
|T| = 1	n 3.2.12 [19]	tn 5.4.2 [31] using 5.3.1 [31]	n^2 3.2.1 [19]	tn^2 4.3.1 [30]
Convex	$\boldsymbol{n}$ 3.2.12 [19]	$\boldsymbol{n}$ **5.3.1**	n^2 3.2.1 [19]	n^2 3.2.2 [17]

Chapter 2

Preliminaries

This chapter is a collection of assumptions, definitions, and lemmas used throughout this dissertation.

2.1 Assumptions

In the following, we summarize the assumptions implicitly used throughout this dissertation.

Notation abuses. We often omit the braces around sets defined extensively, especially for pairs of segments which are to be understood as unordered pairs, i.e., sets of exactly two segments.

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 dissertation, and unless mentioned otherwise, we assume general position of the endpoints of the segments we consider, meaning that any three endpoints are assumed to not be colinear. We also assume that the two endpoints of a segment are distinct. The only time we drop these two assumptions is in the proof of Theorem 3.1.3 [19, 82]. In the proofs of Theorem 3.1.4 [19], Theorem 5.8.1 [33], Theorem 6.2.1 [31], and of Theorem 7.1.1 [31], we extend our standard general position assumptions to also exclude pairs of endpoints with the same y-coordinate.

2.2 Definitions

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

Segment types. Given a multiset of segments S and two (possibly equal) sets P_1, P_2 of endpoints, we say that a segment is a P_1P_2-*segment* if it belongs to S, if one endpoint is in P_1, and if 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 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).

General position implies that two intersecting segments may either cross or share at least one endpoint. To fully remove any self-intersection of a tour without any assumption, other types of intersecting pairs of segments have been considered in [82]. We discuss this matter in Section 3.1.1.

Flips. A flip f in the Π version removing the crossing segments s_1, s_2 and inserting the non-crossing segments s'_1, s'_2 with the same four endpoints is defined as the function that maps any set of segments S satisfying Π, containing s_1 and s_2, and such that $S \cup \{s'_1, s'_2\} \setminus \{s_1, s_2\}$ also satisfies Π, to $f(S) = S \cup \{s'_1, s'_2\} \setminus \{s_1, s_2\}$.

Parameters. The main parameter in all the bounds in this dissertation is n, defined as $|S|$, the number of segments (counted with multiplicity) in the multiset of segments S to be untangled. In the following, we define another parameter used in this dissertation.

Any point set P may be partitioned into $P = C \cup T$ where C is in convex position even if it means that C or T is empty. In this context, we define the parameter t as the number of segments with at least one endpoint in T (the segments with two endpoints in T are counted twice). The parameter t may equivalently be defined as the sum of the degrees of the points in T. In this dissertation, the bounds involving t get better when t is small. It is therefore best to choose the partition $P = C \cup T$ so as to minimize t (this choice may be different from choosing C to be the points defining the convex hull of P).

We also provide fine grain parameterized bounds using the following parameters: $|P|$, $|C|$, $|T|$.

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. Next, we give names to the properties Π which are considered in this dissertation. We classify these properties into the following three types.

1. The *point set properties* are the properties about the position of the endpoints (note that they are also preserved by any insertion choice at every flip).

2. The *degree properties* are constraints on the degree of each point (note that they are preserved by any insertion choice at every flip).

3. The *insertion properties* consist of the empty (i.e., always true) property and of the properties that impose which pair is inserted at *every*[1] flip.

The **point set properties** (also preserved by any insertion choice) considered is this dissertation are the following (see Figure 1.4 or the beginning of Section 1.6):

- `General`: the empty property, i.e., the property which is always true (having a general point set is a given). This notation is only used in the tables of Section 1.6.
- $\texttt{C} \cup \texttt{T}$: the empty property (any point set P may be partitioned into $P = C \cup T$ where C is in convex position even if it means that C or T is empty). This notation is only used in the tables of Section 1.6 to indicate that the bounds are parameterized by t, the sum of the degrees of the points in T.
- `Convex`: having a point set in convex position.
- `Permutation`: having a bipartite point set in convex position such that the red points are separated from the blue points by a line.
- `Redonaline`: having a bipartite point set with red points on the x-axis and blue points above.
- $|\texttt{T}| = \texttt{k}$: having a point set partitioned as $P = C \cup T$ with C in convex position and $|T| = k$.
- `Inout`: having a $|\texttt{T}| = 2$ point set with the one point of T inside the convex hull of C, and the other point of T outside.
- `Inin`: having a $|\texttt{T}| = 2$ point set with the two points of T inside the convex hull of C.
- `Outout`: having a $|\texttt{T}| = 2$ point set with the two points of T outside the convex hull of C.
- `Allout`: having a point set partitioned as $P = C \cup T$ with C in convex position and all the points of T outside the convex hull of C.
- `Separated`: having an `Allout` point set partitioned as $P = C \cup T$ with C in convex position and separated from T by two parallel lines.

The only **degree property** (preserved by any insertion choice) considered is this dissertation is the following:

- `Matching`: being a matching.

The **insertion properties** defining which pair of segments to insert for every flip considered is this dissertation are the following:

- `Multigraph`: the empty property (being a multigraph is a given).

[1]In particular, we do not consider properties such as “admitting a given coloring” where the given coloring has more than two colors.

- `Bipartite`: being bipartite for a given partition of the point set.
- `Cycle`: being a cycle, i.e., being a tour.
- `Tree`: being a tree.

All these properties may be combined as in `Convex Bipartite Matching`; spaces stands for conjunctions and we may omit the empty property or a property that is inferred by the rest.

The flip graph and the choice graph. Consider a property $\Pi = \Pi_1\ \Pi_2\ \Pi_3$ being the conjunction of a point set property Π_1, a degree property Π_2, and an insertion property Π_3. Given a point set P satisfying Π_1 and a function $d^\circ : P \to \mathbb{N}^*$ satisfying Π_2, the *flip graph* of (P, d°, Π_3) is the directed simple graph defined as follows. The set of vertices $V^{\mathtt{R}}$ consists of all the multigraphs satisfying Π_3 with vertices P and a degree function equal to d°. There is a directed arc from a multigraph $S_1 \in V^{\mathtt{R}}$ to a multigraph $S_2 \in V^{\mathtt{R}}$ if there exists a flip in the Π_3 version transforming S_1 into S_2.

The *choice graph* of (P, d°, Π_3) is the directed simple graph constructed from the flip graph of (P, d°, Π_3) by decomposing each arc into two arcs with an extra vertex representing a multigraph after removal and before insertion. Specifically, given a class K of multigraphs satisfying Π, let $\widehat{K}$ be the class of all the multigraphs obtained from the multigraphs of K by applying only the removal part of the flips in the Π_3 version. The vertices of the choice graph are $V^{\mathtt{R}} \cup V^{\mathtt{I}}$ where $V^{\mathtt{I}} = \widehat{V^{\mathtt{R}}}$. For each directed arc from a multigraph $S_1 \in V^{\mathtt{R}}$ to a multigraph $S_2 \in V^{\mathtt{R}}$ in the flip graph, there are two directed arcs in the choice graph: the directed arc from S_1 to $S_1' \in V^{\mathtt{I}}$ where S_1' is the multigraph obtained from S_1 by applying only the removal part of the flip, and the directed arc from S_1' to S_2.

The flip graphs (respectively the choice graphs) in the Π version are the flip graphs (respectively the choice graphs) of (P, d°, Π_3) for all possible P satisfying Π_1 and all possible d° satisfying Π_2.

Strategies. We define a *removal strategy* in the Π version to reach a class K of multigraphs as a function mapping each choice graph in the Π version and each vertex S in $V^{\mathtt{R}} \setminus K$ of this choice graph to one of the arcs from S.

We define an *insertion strategy* in the Π version to reach a class K of multigraphs as a function mapping each choice graph in the Π version and each vertex S in $V^{\mathtt{I}} \setminus \widehat{K}$ of this choice graph to one of the arcs from S. Note that there exists exactly one insertion strategy if Π_3 is not empty.

A *strategy for both removal and insertion choices* in the Π version to reach a class K consists of a removal strategy and an insertion strategy in the Π version to reach a class K.

The class K always includes the class of crossing-free multigraphs satisfying Π; it is most of the time taken to be equal.

Untangle sequences. A sequence of multisets $S_0, \ldots, S_k$ is called a *flip sequence* (in the Π version) and k is called its *length* if S_0 satisfies Π and if for each $i \in \{1, \ldots, k\}$, there exists a flip f_i (in the Π version) such that $S_i = fi(S_{i-1})$. In other words, a flip sequence is a path in the flip graph.

A flip sequence (in the Π version) $S_0, \ldots, S_k$ is called an *untangle sequence* (in the Π version) if S_k is crossing free.

Note that a multiset S_0 satisfying Π, a removal strategy in the Π version to reach a class K of multigraphs, and, if Π is preserved by any the insertion choices, an insertion strategy in the Π version to reach a class K of multigraphs determine a unique flip sequence, and the final multigraph is in K.

The unknown of this dissertation. Consider a property $\Pi = \Pi_1\ \Pi_2\ \Pi_3$ being the conjunction of a point set property Π_1, a degree property Π_2, and an insertion property Π_3. Let n be a non-negative integer. In the following four definitions, $k(S, \mathtt{r}, \mathtt{i})$ denotes the length of the untangle sequence starting at the multiset S of n segments (satisfying Π) and determined by the removal strategy $\mathtt{r}$ (in the Π version to reach the class of crossing-free multigraphs) and by the insertion strategy $\mathtt{i}$ (in the Π version to reach the class of crossing-free multigraphs).

$$\mathbf{d}^{\emptyset}_{\Pi}(n) = \max_{S} \max_{\mathtt{r}} \max_{\mathtt{i}}\ k(S, \mathtt{r}, \mathtt{i})$$

$$\mathbf{d}^{\mathtt{R}}_{\Pi}(n) = \max_{S} \min_{\mathtt{r}} \max_{\mathtt{i}}\ k(S, \mathtt{r}, \mathtt{i})$$

$$\mathbf{d}^{\mathtt{I}}_{\Pi}(n) = \max_{S} \max_{\mathtt{r}} \min_{\mathtt{i}}\ k(S, \mathtt{r}, \mathtt{i}) \qquad \text{(defined only if } \Pi_3 \text{ is empty)}$$

$$\mathbf{d}^{\mathtt{RI}}_{\Pi}(n) = \max_{S} \min_{\mathtt{r}} \min_{\mathtt{i}}\ k(S, \mathtt{r}, \mathtt{i}) \qquad \text{(defined only if } \Pi_3 \text{ is empty)}$$

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 choice graph as a board. The adversary chooses the starting vertex in $V^{\mathtt{R}}$. A game builds a path in the choice graph arc by arc until a crossing-free vertex is reached. The oracle aims at minimizing the length of the path while the adversary aims at maximizing this length. The four definitions corresponds to the length of this resulting path with different choices for the oracle (specified in the exponent). This length is $\mathbf{d}^{\emptyset}$ if the adversary performs all the choices to be done, it is $\mathbf{d}^{\mathtt{R}}$ if the oracle performs all the removal choices while the adversary performs all the remaining choices, it is $\mathbf{d}^{\mathtt{I}}$ if the oracle performs all the insertion choices while the adversary performs all the remaining choices, it is $\mathbf{d}^{\mathtt{RI}}$ if the oracle performs all the removal and all the insertion choices (the adversary only get to choose the starting vertex).

Notice that we define $\mathbf{d}^{\mathtt{I}}_{\Pi}$ and $\mathbf{d}^{\mathtt{RI}}_{\Pi}$ only if Π_3 is empty. Indeed, it is important for the proof of Lemma 2.3.2 that *all* the occurrences of the choices specified in the exponent are performed by the oracle. This is also the reason why we do not consider properties such as "admitting a given coloring" where the given coloring has more than two colors (preserving such a property sets the insertion choice of only certain flips).

2.3 Reductions

In this section, we prove general inequalities between the different versions of $\mathbf{d}$.

Lemma 2.3.1 (choice reduction)**.** *The following inequalities hold for any non-negative integer n, and for any property Π.*

$$\mathbf{d}_{\Pi}^{\mathtt{RI}}(n) \leq \begin{Bmatrix} \mathbf{d}_{\Pi}^{\mathtt{R}}(n) \\ \mathbf{d}_{\Pi}^{\mathtt{I}}(n) \end{Bmatrix} \leq \mathbf{d}_{\Pi}^{\emptyset}(n)$$

Proof. The more choice untangle sequences have, the shorter, hence the inequalities. □

Lemma 2.3.2 (property reduction)**.** *The following inequality holds for any non-negative integer n, for any two properties Π, Π' such that $\Pi \implies \Pi'$, and for any Choices $\in \{\emptyset, \mathtt{R}, \mathtt{I}, \mathtt{RI}\}$.*

$$\mathbf{d}_{\Pi}^{\mathit{Choices}}(n) \leq \mathbf{d}_{\Pi'}^{\mathit{Choices}}(n)$$

Proof. First notice that any untangle sequence in the Π version is also a untangle sequence in the Π' version.

We show that going from $\mathbf{d}_{\Pi}^{\mathit{Choices}}(n)$ to $\mathbf{d}_{\Pi'}^{\mathit{Choices}}(n)$ gives more choice to the adversary while leaving the oracle with essentially the same choice. The adversary may prevent the oracle to benefit from the untangle sequences in the Π' version which are no in the Π version by choosing the starting multiset of segment. In the case where *Choices* $\in \{\emptyset, \mathtt{R}\}$, some insertion choices which were set to preserve property Π may not be set to preserve property Π' anymore. Let us call this phenomenon *insertion choice relaxation.* Therefore, the adversary may also gains choice from insertion choice relaxation. Note that, by definition of $\mathbf{d}^{\mathit{Choices}}$ and in contrast with the adversary, the oracle does not gain any choice because of insertion choice relaxation.

Thus, the inequality holds for any *Choices* $\in \{\emptyset, \mathtt{R}, \mathtt{I}, \mathtt{RI}\}$. □

Lemma 2.3.3 (transfer reduction)**.** *The following inequalities hold for any non-negative integer n, and for any property Π such that preserving Π does not set insertion choice.*

$$\mathbf{d}_{\Pi\ \mathtt{Matching}}^{\mathtt{RI}}(n) \leq \mathbf{d}_{\Pi\ \mathtt{Bipartite\ Matching}}^{\mathtt{R}}(n)$$
$$\mathbf{d}_{\Pi\ \mathtt{Matching}}^{\mathtt{I}}(n) \leq \mathbf{d}_{\Pi\ \mathtt{Bipartite\ Matching}}^{\emptyset}(n)$$

Proof. Going from $\mathbf{d}_{\Pi\ \mathtt{Matching}}^{\mathtt{RI}}(n)$ (respectively $\mathbf{d}_{\Pi\ \mathtt{Matching}}^{\mathtt{I}}(n)$) to $\mathbf{d}_{\Pi\ \mathtt{Bipartite\ Matching}}^{\mathtt{R}}(n)$ (respectively $\mathbf{d}_{\Pi\ \mathtt{Bipartite\ Matching}}^{\emptyset}(n)$), we transfer insertion choice from the oracle to preserving the `Bipartite` property. Any matching being 2-colorable, the adversary keeps the same choice in the starting matching (respectively in the starting matching and in removal choices), yielding the two inequalities. □

2.4 Triangle Argument

The following simple fact, first stated informally then as a lemma without proof, is used in many of the proofs presented in this dissertation. We refer to it (or its contraposition) as the *triangle argument* or implicitly use it.

"If a line enters a triangle through one edge, then it must leave through another edge."

Lemma 2.4.1 (triangle argument)**.** *If a line ℓ crosses the edge pq of a (possibly flat) triangle pqo, then ℓ intersects $(po \cup qo) \setminus \{p, q\}$, the union of the two other edges excluding p and q, at exactly one point.*

2.5 Potential Argument

In this section, we state another simple lemma which is also used in many proofs throughout this dissertation.

Lemma 2.5.1 (potential argument). *Let $\varphi : I \to J$ be a function. If J is a finite set and if, for all $j \in J$, $\varphi^{-1}(\{j\})$ is a finite set, then I is a finite set and we have the following identity.*

$$|I| = \sum_{j \in J} \left|\varphi^{-1}(\{j\})\right|$$

In this context, we call the function φ a *potential*. Our most frequent use of Lemma 2.5.1 is when J is an integer interval and φ is decreasing to upper bound the length $|I|$ of a sequence (e.g., a flip sequence) $(S_i)_{i \in I}$ by $|J|$. Intuitively, the potential of a step of a process intend to measure how far this step is from the end of the process (e.g., from being crossing free), and a decrease of potential intend to measure the progress made by each step of the process (e.g., each flip).

2.6 Splitting Lemma

The notion of splitting first appears in [19] (see the proof of Theorem 3.2.12 [19]), then in [33], and in [31] where the notion is generalized. In the following, we first give this generalized definition of splitting, then state the corresponding lemma, and finally give some examples.

Consider a finite point set P decomposed into $P = \bigcup_{i=1}^{k} P_i$ (this decomposition is not necessarily a partition). We say that the decomposition $\bigcup_{i=1}^{k} P_i$ *splits* the point set P if no segment of $\binom{P_i}{2}$ crosses a segment of $\binom{P_j}{2}$ for all $i \neq j$. Now, consider a finite multiset S of segments *partitioned* into $S = \bigcup_{i=1}^{k} S_i$ and let P be the set of endpoints of S and P_i be the set of endpoints of S_i. We say that the partition $\bigcup_{i=1}^{k} S_i$ *splits* the multiset S if the decomposition $\bigcup_{i=1}^{k} P_i$ *splits* the point set P (Figure 2.1). In the `Bipartite` version where the point set is partitioned into red and blue points, we refine these definitions as follows. We say that the decomposition $\bigcup_{i=1}^{k} P_i$ *splits* the point set P if no segment of $\binom{P_i}{2}$ with one red endpoint and one blue endpoint crosses a segment of $\binom{P_j}{2}$ with one red endpoint and one blue endpoint for all $i \neq j$ (Figure 2.1(e)).[2]

The following two facts imply Lemma 2.6.1 sated next.

- The four endpoints of the removed and inserted segments of any flip of S all belong to the same P_i for some i.
- If f_i is a flip of S_i and f_j is a flip of S_j, then they *commute*, i.e., $f_i(f_j(S)) = f_j(f_i(S))$.

Lemma 2.6.1 (splitting; [19, 31]). *Consider a partition $S = \bigcup_{i=1}^{k} S_i$ splitting a finite multiset S of segments.*

[2] The splitting definition may be refine even further in any version by considering in each $\binom{P_i}{2}$ only the segments which are inserted in at least one flip sequence of S.

A sequence of flips applied to S is partitioned into k subsequences, the i-th subsequence consisting of all the flips which apply to S_i, and these subsequences commute *in the sens that any flip in one subsequence commutes with a flip in another subsequence.*

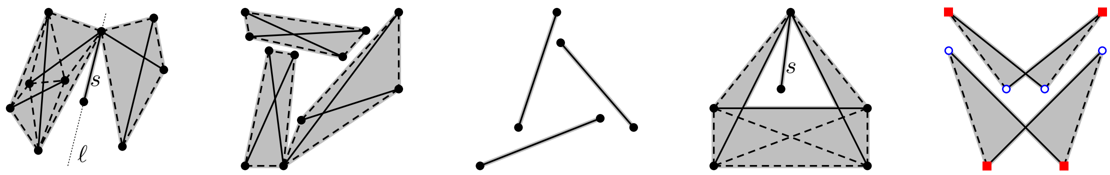

(a) (b) (c) (d) (e)

Figure 2.1: Examples of splitting partitions of multisets of segments where each partition is shaded and where the segments inserted in some flip sequence are dashed. Each occurrence of the segment s is uncrossable. (a) The line ℓ splits the segments in three partitions, one of which consisting of the segment s alone. (b) A partition which is not induced by any line splitting the segments in three. (c) Crossing-free segments, each of them forming its own partition for splitting. (d) A partition splitting the segments in two. One partition consists of the segment s alone (e) A partition segments in the `Bipartite` version.

Splitting (Lemma 2.6.1 [19, 31]) is used throughout this dissertation for inductive arguments. Some important examples of splitting follow.

- If a line ℓ does not cross any segment of a finite multiset of segment S, then the partition $S = S_0 \cup S_1 \cup S_2$ where S_0 are the segments of S which lie are contained in ℓ, where S_1 are the segments of S which lie on one side of ℓ, and where S_2 are the segments of S which lie on the other side of ℓ splits S. In this case, we say that ℓ *splits* S (Figure 2.1(a)).

- More generally, if there exists a decomposition $P = \bigcup_{i=1}^{k} P_i$ of the point set P induced by a partition $S = \bigcup_{i=1}^{k} S_i$ of the multiset of segments S such that the convex hulls of the P_i have pairwise disjoint interiors, then the partition splits S. In this case, we say that the splitting is *convex* (Figure 2.1(b)).

- In particular, in such a convex splitting, some of the convex hulls may consist of only one segment, as it is the case if S is crossing free (Figure 2.1(c)).

- However, some splitting partitions are not convex; Figure 2.1(d) shows a multiset of segments S and a segment $s \in S$ such that the partition $\{s\} \cup S \setminus \{s\}$ splits S while not being convex.

- If a singleton $\{s\}$ is one of the partitions in a splitting partition, then we say that the segment s is *uncrossable* (the segment s is uncrossable in Figure 2.1(a) and Figure 2.1(d), and also all the three segments of Figure 2.1(c)).[3]

[3]This definition of uncrossable generalizes the one in [31].

Chapter 3

Literature Review

In this chapter, we present the results from the literature falling in problematic of this dissertation, namely from [17, 19, 72, 82, 85]. We omit the proof of Theorem 3.2.9 and give only a structure for the proofs of Theorem 3.2.4 and Theorem 3.1.5. We adapt and generalize the presentation of the rest of the proofs. First, we give a short review of each of these five articles.

***Untangling a Traveling Salesman Tour in the Plane* by Jan Van Leeuwen and Anneke A. Schoone in 1980 [82, 83]** This article[1] is the seminal work of all the following articles and of the present dissertation. First, they provide a formal proof of the folklore theorem which asserts that the shortest tours have no self intersection if the cities are not all colinear. Last but not least, they prove that any untangle sequence is of length at most n^3 (n being the number of segments in the tours), more formally that:

- $\mathbf{d}^{\emptyset}_{\texttt{Cycle}}(n) \leq n^3$ (Theorem 3.1.3).

This proof is written for the `Cycle` version (and do not assume general position for the point set) but it is straightforward to adapt to the `Matching` version [19][2], or to generalize it to the `Multigraph` version (we present this generalisation in Theorem 3.1.3). In this proof, they introduce a potential function (the function mapping a tour to the number of intersections between a set of lines and the set of segments of the tour). This idea is at the core of many results in [17, 19, 30, 31, 33] and therefore in the present dissertation.

***The Number of Flips Required to Obtain Non-crossing Convex Cycles* by Yoshiaki Oda and Mamoru Watanabe in 2007 [72].** This article provides the following upper and lower bounds:

- $\mathbf{d}^{\mathtt{R}}_{\texttt{Convex Cycle}}(n) \leq 2n - 7$ for $n \geq 4$ (Theorem 3.2.7),
- $n - 2 \leq \mathbf{d}^{\mathtt{R}}_{\texttt{Convex Cycle}}(n)$ for $n = 5$ and for $n \geq 7$.

[1] The technical report [82] also appeared in a workshop [83].

[2] However, the paper [82] is not cited in [19], hinting that they perhaps rediscovered the result independently.

The article also studies $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Cycle}}$ restricted to some specific classes of convex polygons. The last section conjectures that $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Cycle}}(n) = n-2$ for $n \geq 7$, then computationally checks this conjecture for $7 \leq n \leq 11$. A list of convex polygons satisfying the conjecture is also provided.

***On the Maximum Switching Number to Obtain Non-crossing Convex Cycles* by Ro–Yu Wu, Jou–Ming Chang, and Jia–Huei Lin in 2009 [85].** This article proves the conjecture of [72], namely:

- $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Cycle}}(n) \leq n-2 \quad$ for $n \geq 7$ (Theorem 3.2.9).

***Flip Distance to a Non-crossing Perfect Matching* by Édouard Bonnet and Tillmann Miltzow in 2016 [19].** This article proves the following bounds:

- $\mathbf{d}^{\emptyset}_{\mathtt{Matching}}(n) \leq n^3$ (Theorem 3.1.3),
- $\binom{n}{2} \leq \mathbf{d}^{\emptyset}_{\mathtt{Convex\ Permutation\ Matching}}(n)$ (Theorem 3.2.1),
- $n-1 \leq \mathbf{d}^{\emptyset}_{\mathtt{Matching}}(n)$ (Theorem 3.2.12),
- $\mathbf{d}^{\mathtt{I}}_{\mathtt{Matching}}(n) \leq \frac{n^2}{2}$ (Theorem 3.1.4).

The proof provided for the third bound is in fact a proof of the following stronger result: $n-1 \leq \mathbf{d}^{\mathtt{RI}}_{\mathtt{Matching}}(n)$. The article additionally contains the following conjecture: $\mathbf{d}^{\emptyset}_{\mathtt{Matching}} = O(n^2)$, which we generalize to the `Multigraph` version in Conjecture 1. This conjecture remains open.

***Flip Distance to some Plane Configurations* by Ahmad Biniaz, Anil Maheshwari, and Michiel Smid in 2019 [17].** This article is the most recent of the five. The following bounds are proven:

- $\mathbf{d}^{\mathtt{I}}_{\mathtt{Matching}}(n) = O(n\sigma(P)) \quad$ where $\sigma(P)$ is the spread of the point set P (Theorem 3.1.5),
- $\mathbf{d}^{\emptyset}_{\mathtt{Convex\ Matching}}(n) \leq \binom{n}{2}$ (Theorem 3.2.2),
- $\mathbf{d}^{\mathtt{RI}}_{\mathtt{Convex\ Matching}}(n) \leq n-1$ (Theorem 3.2.13),
- $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Bipartite\ Matching}}(n) \leq 2n-2$ (Theorem 3.2.10),
- $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Tree}}(n) = O(n \log n)$ (Theorem 3.2.11),
- $\mathbf{d}^{\mathtt{R}}_{\mathtt{Redonaline\ Matching}}(n) \leq n(n-1)$ (Theorem 3.3.1).

The proofs of the two first bounds in fact holds for the `Convex Multigraph` version. The proofs in [17] of second and the third bounds are both very similar to the proof of the upper bound on $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Cycle}}$ in [72], and all use the same notion of *depth* of a segment. We present proofs based on a similar but distinct notion, which we call the *crossing depth* of a segment.

3.1 General Point Sets

In this section, we prove the literature results concerning general point sets.

3.1.1 Untangling with No Choice

In this section, we prove the only literature bound on tho length of untangle sequences with no choice: the unchallenged 1980 cubic upper bound on $\mathbf{d}^{\emptyset}_{\texttt{Multigraph}}$. The literature contains no lower bound specific to untangle sequences with no choice, the best lower bound known on $\mathbf{d}^{\emptyset}_{\Pi}$ for any Π is discussed next.

Lower Bound

To our knowledge, the longest untangle sequences contained in the literature are all with point sets in convex position and their length is $\Theta(n^2)$ (Theorem 3.2.1). This could lead to conjecture that some of the longest untangle sequences are indeed attained with point sets in convex position. We disprove this conjecture in Theorem 4.2.1. However, the following conjecture, formulated for the `Matching` version in [19], is not disproved and matches our current belief.

Conjecture 1 ([19])**.** *In the* `Multigraph` *version, any untangle sequence is of length at most* $O(n^2)$*.*

In other words, we have the following upper bound: $\mathbf{d}^{\emptyset}_{\texttt{Multigraph}} = O(n^2)$*.*

Upper Bound

In the 1980 technical report [82] and the year after in the workshop [83], Jan van Leeuwen and Anneke A. Schoone proved Theorem 3.1.3 in the `Cycle` version. It is the seminal result of the present dissertation. An adaptation of the proof to the `Bipartite Matching` version has been presented in [19]. We present here a generalization to the `Multigraph` version.

It is notable that, in [82], general position is not assumed. In the following, we also present the proof of Theorem 3.1.3 without any assumption on the point set. This leads to consider a generalization of flips defined as follows. A *general flip* in the Π version is an operation removing two intersecting segments from a multiset of segments and inserting two non-intersecting segments with the same four endpoints while preserving a property Π of the multiset of segments. In this setting, segments degenerated to a point are allowed; they are modeled by loops in the corresponding multigraph. There are seven *types* of general flips depending on the *type* of intersection of the removed segments (Figure 3.1).

- *Type I*: the removed segments are crossing (the general flips of type I are the flips).
- *Type II*: the removed segments intersect at exactly one point which is one of the endpoints of exactly one of the removed segments.
- *Type III*: the removed segments intersect at exactly one point which is one of the endpoints of both removed segments.

- *Type IV*: the removed segments intersect at infinitely many points, their four endpoints are distinct, and one removed segment contains the other.
- *Type V*: the removed segments intersect at infinitely many points, their four endpoints are distinct, and none of the removed segments contains the other.
- *Type VI*: the removed segments intersect at infinitely many points, and they share exactly one endpoint.
- *Type VII*: the removed segments intersect at infinitely many points, and their are identical.

Note that in the general flips of type III to VII, the insertion choice is set by the condition that the inserted segments do not intersect. If, for a given pair of segments, this insertion choice does not preserves the property Π, then the pair of segments cannot be flipped.

We prove Theorem 3.1.3 using Lemma 3.1.1 and Lemma 3.1.2 which we state and prove next. Lemma 3.1.1 relies on a third elementary lemma (Lemma 2.4.1). The following definitions are used.

We define $\Lambda_\ell(S)$ as the number of segments of the multiset S crossing the line ℓ (counted with their multiplicity). If L is a set (or a multiset) of lines, we define $\Lambda_L(S)$ as the sum of the $\Lambda_\ell(S)$ for $\ell \in L$.

Lemma 3.1.1. *Consider a line ℓ and a multiset of segments S. A general flip of S does not increase $\Lambda_\ell(S)$.*

Proof. Let p_1p_2 and p_3p_4 be the two segments removed by the general flip, let o be one of their intersection points, and let p_1p_3 and p_2p_4 be the two segments inserted by the general flip. We show that any crossing pair consisting of ℓ and an inserted segment is reflected by a crossing pair consisting of ℓ and a removed segment.

If both removed segments cross ℓ, or if none of the inserted segments cross ℓ, then the lemma holds. The only remaining case is when ℓ crosses only one inserted segment, say p_1p_3. Then, ℓ intersects the (possibly flat) triangle p_1p_3o. Thus by the triangle argument (Lemma 2.4.1), ℓ also intersects $p_1o \cup p_3o$ at exactly one point o' (which is not p_1 or p_3). This intersection o' is not o either, as we have assumed that at least one of the removed segments does not cross ℓ. Therefore, o' is not p_2 or p_4 and the lemma holds. □

Before we state and prove Lemma 3.1.2, we need the following additional definitions. Given two distinct points p_1, p_2 of a point set P, let $p_1p_2^+$ and $p_1p_2^-$ by two lines crossing the segment p_1p_2 and being as close to the line p_1p_2 as necessary to satisfy the following properties for all segments s with endpoints in P[3].

- If the endpoints of s do not lie on the line p_1p_2, then $p_1p_2^+$ (respectively $p_1p_2^-$) intersects s if and only if the line p_1p_2 intersects s.
- If exactly one of the endpoints of s lies on the line p_1p_2, then exactly one of $p_1p_2^+$ or $p_1p_2^-$ intersects s.

[3]In [19], lines parallel to p_1p_2 are used instead. This choice does not work well if general position is not assumed.

In other words, the two partitions of $P \setminus p_1p_2$ induced by the lines $p_1p_2^+$ and $p_1p_2^-$ are the same as the partition induced by p_1p_2 and are different from each other. We say that $p_1p_2^+$ and $p_1p_2^-$ are the *witnesses* of the segment p_1p_2. Let L be the multiset of the lines $p_1p_2^+$ and $p_1p_2^-$ for all pairs of points $p_1, p_2 \in P$[4].

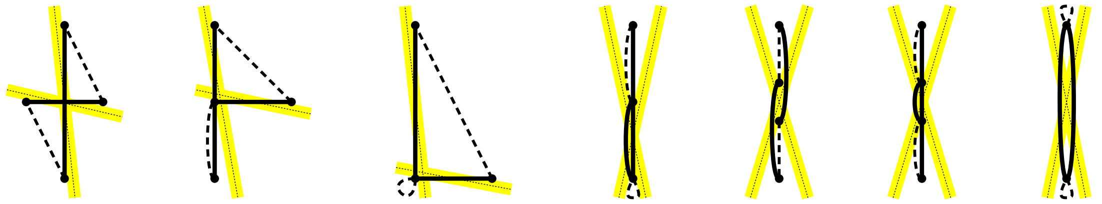

(I) (II) (III) (IV) (V) (VI) (VII)

Figure 3.1: The seven types of general flips (removed segments are plain bold, inserted segments are dashed) with the two witness lines (doted and highlighted) in L decreasing Λ_ℓ.

Lemma 3.1.2. *For any general flip, there exists at least two lines $\ell \in L$ such that Λ_ℓ decreases by 2.*

Proof. An exhaustive case analysis of the seven types of general flips yields that each removed segment has a witness $\ell \in L$ such that Λ_ℓ decreases by 2 (Figure 3.1). □

Theorem 3.1.3 ([19, 82]). *Consider a multiset S of n segments with endpoints P.*

In the `Multigraph` *version, any untangle sequence of S is of length at most $\frac{1}{2}n\binom{|P|}{2} = O(n|P|^2) = O(n^3)$.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\emptyset}_{\texttt{Multigraph}}(n) \leq \frac{1}{2}n\binom{|P|}{2} = O(n|P|^2) = O(n^3).$$

Proof. A consequence of Lemma 3.1.1 and Lemma 3.1.2 is that Λ_ℓ decreases by at least 4 at each general flip. Moreover, each of the $2\binom{|P|}{2}$ lines in L crosses at most n segments. Thus, the function Λ_ℓ takes integer values between 0 and $2n\binom{|P|}{2}$, hence the theorem. □

3.1.2 Untangling with Insertion Choice

In this section, we present the two upper bounds on $\mathbf{d}^{\mathtt{I}}_{\texttt{Multigraph}}$ from the literature: a general upper bound and an upper bound parameterized by the spread of the point set (defined next). These results do not carry over to the `Cycle` version where insertions are set to preserve the `Cycle` property.

The literature contains no lower bound specific to insertion choice, the best lower bound known on $\mathbf{d}^{\mathtt{I}}$ is a corollary of Theorem 3.2.12.

[4]In [82], the set of the lines p_1p_2 for all pairs of points $p_1, p_2 \in P$ is used instead. This choice leads to a more tedious analysis if general position is not assumed. However, if general position is assumed, the analysis is smooth and the upper bound slightly improves to $\frac{1}{2}n\left(\binom{|P|}{2} - 2|P| - 3\right)$ (a given segment does not cross the $2|P| - 3$ lines containing at least one of its endpoints). Moreover, in the `Bipartite Matching` version, the lines defined by a red point and a blue point suffice, further improving the upper bound to $\binom{n}{2}(n-1)$.

General Upper Bound

The following theorem provides a quadratic upper bound and relies on insertion choice. We extend our standard general position assumptions to also exclude pairs of endpoints with the same y-coordinate.

Theorem 3.1.4 ([19]). *Consider a multiset S of n segments with endpoints P.*

In the `Multigraph` *version, there exists an insertion strategy such that for any removal strategy, the resulting untangle sequence of S is of length at most $\frac{n}{2}(|P|-2) = O(n\,|P|) = O(n^2)$.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\mathtt{I}}_{\mathtt{Multigraph}}(n) \leq \frac{n}{2}(|P|-2) = O(n\,|P|) = O(n^2).$$

Proof. Let $p_1, \ldots, p_{|P|}$ be the points of P sorted vertically from top to bottom. Consider a flip involving the points p_a, p_b, p_c, p_d with $a < b < c < d$. There are two cases whether the removed pair of segments is $p_a p_c, p_b p_d$ or $p_a p_d, p_b p_c$. In both cases, we choose to insert the pair of segments $p_a p_b, p_c p_d$.

We now prove that this insertion choice leads to the claimed upper bound. 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 $\eta(S)$ as the sum of $\eta(p_i p_j)$ for $p_i p_j \in S$. The function η is equal to Λ_L where L is a set of $|P| - 1$ horizontal lines that partition the plane with one point of P per partition.

Notice that $n \leq \eta(S) \leq n(|P| - 1)$. It is easy to verify that, in both cases for the removed pair of segments, a flip decreases $\eta(S)$ by at least 2. Hence, the number of flips is $\frac{n}{2}(|P| - 2) = O(n\,|P|)$. □

Parameterized Upper Bound

Next, we state Theorem 3.1.5; the statement is identical to Theorem 3.1.4 except that the upper bound is $O(n\sigma(P))$ instead of $O(n^2)$. This calls for the following definition. If P is a set of points, let $\sigma(P)$ be the *spread* of the point set P, that is, the ratio between the maximum and minimum distances among pairs of points in P.

Note that the spread $\sigma(P)$ is not invariant by affine transformation, and thus is not an invariant of the problem. Recall that the problem of untangling a multigraph only depends on the the order type of the point set[5]. Therefore, Theorem 3.1.5 still holds when replacing $\sigma(P)$ by the minimum of $\sigma(P')$ P' varies over all the point sets of the same order type as P.

Also note that Theorem 3.1.5 improves on Theorem 3.1.4 only when $\sigma(P) = o(n)$. As the spread $\sigma(P)$ of a point set P is always lower bounded by $\Omega(\sqrt{n})$(see e.g., [79]), the best improvement is attained for *dense* point sets P, where $\sigma(P) = \Theta(\sqrt{n})$. Dense point sets appear in practical situations and have been studied in several contexts (see e.g., [39]).

[5]An alternative would be the *intrinsic spread* defined as the ratio between the maximum and the minimum area among triplets of points in P proposed in [48]. The intrinsic spread is invariant by affine transformation, but is still not an invariant of the problem.

Theorem 3.1.5 ([17]). *Consider a multiset S of n segments with endpoints P.*

In the `Multigraph` *version, there exists an insertion strategy such that for any removal strategy, the resulting untangle sequence of S is of length $O(n\sigma(P))$.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\mathtt{I}}_{\mathtt{Multigraph}}(n) = O(n\sigma(P)).$$

The proof of Theorem 3.1.5 relies on Lemma 3.1.8, itself relying on Lemma 3.1.6 and Lemma 3.1.7. As the three lemmas are only technical lemmas, we only prove Theorem 3.1.5 assuming Lemma 3.1.8.

Lemma 3.1.6 ([17]). *Let p_1p_2 and p_3p_4 be two crossing segments. Let o be their crossing point. Let μ be the minimum distance between any pair of points in $\{p_1, p_2, p_3, p_4\}$. If the angle $\angle p_3op_2 \leq \frac{\pi}{3}$, then $\mu \leq (p_1p_2 + p_3p_4) - (p_2p_4 + p_3p_2)$.*

Lemma 3.1.7 ([17]). *Let p_1p_2 and p_3p_4 be two crossing segments. Let o be their crossing point. Let μ be the minimum distance between any pair of points in $\{p_1, p_2, p_3, p_4\}$. If the angle $\angle p_3op_2 = \frac{\pi}{2}$, then $\frac{2-\sqrt{2}}{2}\mu \leq (p_1p_2 + p_3p_4) - (p_2p_4 + p_3p_2)$.*

Lemma 3.1.8 ([17]). *Let p_1p_2 and p_3p_4 be two crossing segments. Let o be their crossing point. Let μ be the minimum distance between any pair of points in $\{p_1, p_2, p_3, p_4\}$. If the angle $\angle p_3op_2 \leq \frac{\pi}{2}$, then $\frac{2-\sqrt{2}}{4}\mu \leq (p_1p_2 + p_3p_4) - (p_2p_4 + p_3p_2)$. In other words, inserting the pair of segments facing the two smallest angles among the four angles defined by the removed segments decreases the total length by at least $\frac{2-\sqrt{2}}{4}\mu$ (the constant $\frac{2-\sqrt{2}}{4}$ is not optimized).*

The idea of the proof of Theorem 3.1.5 is to use insertion choice to systematically insert a pair of segments such that the sum of the lengths of all the segments in S decreases by a constant factor.

Proof of Theorem 3.1.5. Let $\mu(P)$ be the minimal distance between two points of P. Let $\Gamma(S)$ be the sum of the lengths of all the segments in S. We have $\Gamma(S) = O(n\mu(P)\sigma(P))$. We systematically insert the pair of segments facing the two smallest angles among the four angles defined by the removed segments. By Lemma 3.1.8, this insertion strategy decreases $\Gamma(S)$ by at least $\frac{2-\sqrt{2}}{4}\mu(P)$, yielding Theorem 3.1.5. □

3.2 Convex Position Point Sets

In this section, we prove the literature results concerning point sets in convex position. In this *convex case*, the geometry of the point set is fully encoded by the cyclic order of the points on the convex hull boundary, leading to fully combinatorial proofs.

3.2.1 Untangling with No Choice

In this section, we prove that $\binom{n}{2}$ is both a lower bound on $\mathbf{d}^{\emptyset}_{\mathtt{Convex\ Permutation\ Matching}}$ and an upper bound on $\mathbf{d}^{\emptyset}_{\mathtt{Convex\ Multigraph}}$.

These two bounds together lead to the following identities.

$$\begin{aligned}\forall \Pi \in \{&\texttt{Multigraph},\\ &\texttt{Bipartite Multigraph},\\ &\texttt{Matching},\\ &\texttt{Bipartite Matching},\\ &\texttt{Redonaline Matching},\\ &\texttt{Permutation Matching}\} \quad \mathbf{d}^{\emptyset}_{\texttt{Convex }\Pi}(n) = \binom{n}{2}\end{aligned}$$

Note that the proof of the lower bound (Theorem 3.2.1) presented next does not hold in the `Cycle` version or in the `Tree` version. However, the reductions presented in Theorem 4.1.1 extend this lower bound asymptotically, yielding $\mathbf{d}^{\emptyset}_{\texttt{Convex }\Pi} = \Theta(n^2)$ for $\Pi \in \{\texttt{Cycle}, \texttt{Tree}\}$.

Lower Bound

In this section, we prove the following lower bound.

Theorem 3.2.1 ([19]). *There exists a set S of n segments forming a bipartite matching with red points on a line ℓ_1 and blue points on a line ℓ_2 parallel to ℓ_1 such that there exist (in the* `Convex Permutation Matching` *version) an untangle sequence of S of length $\binom{n}{2}$.*

In other words, we have the following lower bound:

$$\mathbf{d}^{\emptyset}_{\texttt{Convex Permutation Matching}}(n) \geq \binom{n}{2} = \Omega(n^2).$$

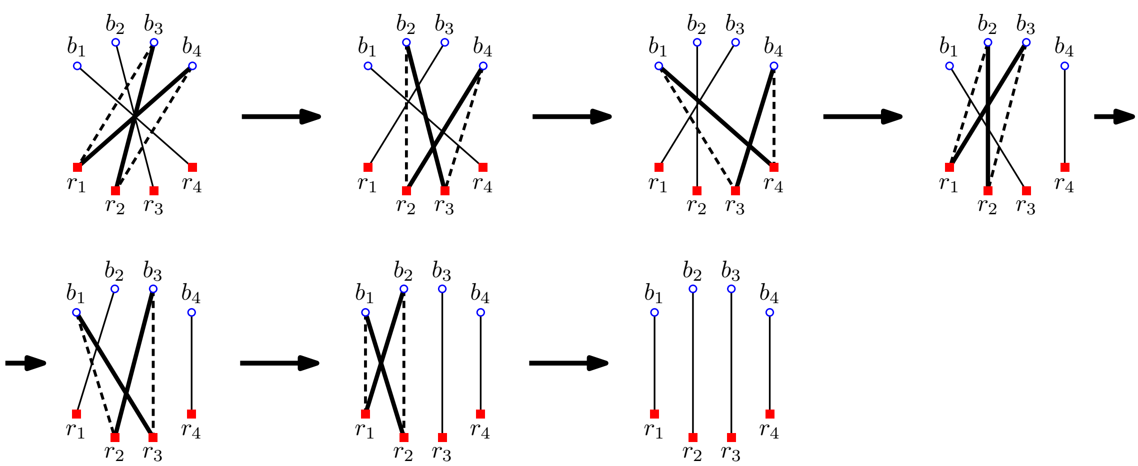

Figure 3.2: A lower-bounding bubble sort untangle sequence used in the proof of Theorem 3.2.1.

Proof. Let $R = \{r_1, \ldots, r_n\}$ be a set of n red points. Let $B = \{b_1, \ldots, b_n\}$ be a set of n blue points. There exists a canonical bijection between $\mathfrak{S}_n$, the permutations of $\{1, \ldots, n\}$, and the bipartite matchings on R, B.

Moreover, if the red and blue points are on the boundary of a convex body, the red points being on one side of a horizontal line and the blue points being on the other side, with red and blue points order from left to right, then the only crossing-free bipartite matching on R, B corresponds to the identity in $\mathfrak{S}_n$, and the crossing pairs of segments in a matching corresponds to the inversions. This hypothesis is satisfied as the red points lie on a line ℓ_1 and the blue points lie on a line ℓ_2 parallel to ℓ_1.

Therefore, a bubble sort of the permutation mapping any $i \in \{1, \ldots, n\}$ to $n+1-i$ corresponds to an untangle sequence with $\binom{n}{2}$ flips (Figure 3.2)). □

Upper Bound

In this section, we prove the following upper bound.

Theorem 3.2.2 ([17]). *Consider a multiset S of n segments with endpoints P in convex position and with m crossing pairs of segments. In the* `Convex Multigraph` *version, all untangle sequences of S have length at most m.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\emptyset}_{\texttt{Convex Multigraph}}(n) \le \binom{n}{2} = O(n^2).$$

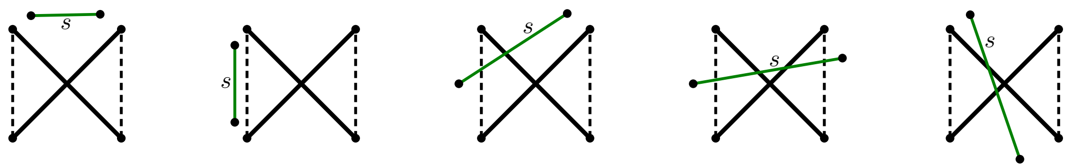

Figure 3.3: The five convex positions of s with respect to the removed and inserted pairs. This is used in the proof of Lemma 3.2.3.

Lemma 3.2.3. *Consider three segments with endpoints in convex position. In the* `Convex Multigraph` *version, flipping two segments does not increase the number of crossing pairs involving the third segment.*

The proof we present hereafter is syntactically identical to the proof of Theorem 3.1.3 (up to replacing the line ℓ by the segment s) which uses a triangle argument (Lemma 2.4.1). The triangle argument normally applies to a line, but a segment behaves like a line if it is long enough, which is the case with a convex point set. In fact, the triangle argument with a long segment instead of a line translates into a combinatorial argument if the endpoints involved are in convex position.

Note that Lemma 3.2.3 also admits a proof by exhaustive case analysis as shown in Figure 3.3. Not only is this exhaustion informative, but it is also used in the proof of Lemma 5.1.4.

Proof of Lemma 3.2.3. Let p_1p_2 and p_3p_4 be the two segments removed by the flip, let o be their intersection point, and let p_1p_3 and p_2p_4 be the two segments inserted by the flip. Let s be the third segment. We show that any crossing pair involving s and an inserted segment is reflected by a crossing pair involving s and a removed segment.

If both removed segments cross s, or if none of the inserted segments cross s, then the lemma holds. The only remaining case is when s crosses only one inserted segment, say p_1p_3. Then, s intersects the (possibly flat) triangle p_1p_3o. Thus by the triangle argument (Lemma 2.4.1), s also intersects $p_1o \cup p_3o$ at exactly one point o' (which is not p_1 or p_3). This intersection o' is not o either, as we have assumed that at least one of the removed segments does not cross s. Therefore, o' is not p_2 or p_4 and the lemma holds. □

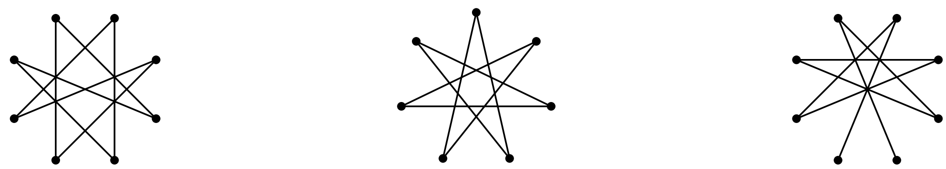

(a) (b) (c)

Figure 3.4: (a) A polygon maximizing the number of crossing pairs of segments with an odd number of segments. (b) A polygon maximizing the number of crossing pairs of segments with an even number of segments. (c) A tree maximizing the number of crossing pairs of segments.

Proof of Theorem 3.2.2. Let $\chi(S)$ be the number of crossing pairs of segments of S. Notice that $\chi(S)$ takes integer values between 0 and $\binom{n}{2}$. Lemma 3.2.3 ensures that χ decreases by at least 1 at a flip; Theorem 3.2.2 follows. □

Note that this upper bound can be refined in the `Convex Cycle` version using the fact that the maximum number of crossing pair of segments in a polygon with n segments is $\binom{n}{2} - n$ if n is odd and $\binom{n}{2} - \frac{3}{2}n + 1$ if n is even [46]. The case where n is odd is attained with a polygon where each segment crosses all the other segments except its neighbors (Figure 3.4(a)). The case where n is even is less obvious (Figure 3.4(b)). Similarly, in the `Convex Tree` version, the maximum number of crossings $\binom{n}{2} - n + 1$ is attained for a path where each segment crosses all the other segments except its neighbors (Figure 3.4(c)).

3.2.2 Untangling with Removal Choice

In this section, we prove linear lower and upper bounds on $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Cycle}}$ (settling even the exact value) and an $O(n \log n)$ upper bound on $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Tree}}(n)$ from the literature.

Lower Bound

In this section, we prove the following lower bound.

Theorem 3.2.4 ([72]). *There exists polygons with n segments and endpoints in convex position such that (in the* `Convex Cycle` *version) any untangle sequence is of length at least $n-3$ if $n \in \{3,4,6\}$ and at least $n-2 = \Omega(n)$ if $n = 5$ or $n \geq 7$.*

In other words, we have the following lower bound:

$$\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Cycle}}(n) \geq \begin{cases} n-3 & \text{for } n \in \{3,4,6\} \\ n-2 = \Omega(n) & \text{for } n = 5 \text{ and } n \geq 7. \end{cases}$$

The proof of relies on the next two lemmas whose statement use the following definitions. The proofs of these two lemmas are omitted.

For odd n, we define $\mathcal{C}_{n,2}$ as the polygon on n points in convex position with edges defining a cycle skipping every other point in cyclic order on the convex hull.

For even n, we define $\mathcal{D}_n$ as the polygon on n points in convex position with edges defining the following cycle described with the indices of the points sorted in counterclockwise order on the convex hull.

$$(1, 3, 5, \ldots, n-5, n-2, n, n-3, n-1, 2, 4, 6, \ldots, n-4)$$

Lemma 3.2.5. *Let n be a positive odd integer at least 5. In the* `Convex Cycle` *version, any untangle sequence of the polygon $\mathcal{C}_{n,2}$ is of length at least $n-2$.*

Lemma 3.2.6. *Let n be a positive even integer at least 8. In the* `Convex Cycle` *version, any untangle sequence of the polygon $\mathcal{D}_n$ is of length at least $n-2$.*

We now provide a short proof of Theorem 3.2.4 assuming Lemma 3.2.5 and Lemma 3.2.6.

Proof of Theorem 3.2.4. It is easy to check that $\mathbf{d}^{\mathrm{R}}_{\texttt{Convex Cycle}}(3) = 0$, $\mathbf{d}^{\mathrm{R}}_{\texttt{Convex Cycle}}(4) = 1$, $\mathbf{d}^{\mathrm{R}}_{\texttt{Convex Cycle}}(5) = 3$, and $\mathbf{d}^{\mathrm{R}}_{\texttt{Convex Cycle}}(6) = 3$. Lemma 3.2.5 and Lemma 3.2.6 yield Theorem 3.2.4 for $n \geq 7$. □

Upper Bounds

In this section, first we define the *crossing depth* (used in the proofs of Theorem 3.2.7, Theorem 3.2.10, Theorem 3.2.13, and Theorem 5.2.1) of a segment. All the proofs using these definition are very similar.

Then, we state four theorems providing linear or quasilinear upper bounds on $\mathbf{d}^{\mathrm{R}}_{\texttt{Convex}}$: two for the `Cycle` version, one for the `Bipartite Matching` version, and one for the `Tree` version. We provide proofs for only two of them. We omit the proof of Theorem 3.2.9 which improves the constants of the upper bound of Theorem 3.2.7 to reach exact tightness. We delay the proof of Theorem 3.2.11 to the proof of the more general Theorem 5.2.1.

Crossing depth. Consider a multiset S of n segments with endpoints P in convex position. Let $p_1, \ldots, p_{|P|}$ be the points of P sorted in counterclockwise order along the convex hull boundary. Consider a segment $p_a p_b$, assuming without loss of generality that $a < b$.

We define the *crossing depth* [31][6] $\delta_\times(p_a p_b)$ as the number of points in $p_{a+1}, \ldots, p_{b-1}$ that are an endpoint of a segment in S that crosses any segment in S (not necessarily $p_a p_b$).

Note that the crossing depth of a segment is a non-negative integer at most $|P| - 2$. Also note that, by splitting (Lemma 2.6.1 [19, 31]), the segments of crossing depth 0 are uncrossable and can no longer participate in a flip.

The following theorem provides a linear upper bound in the `Convex Cycle` version.

[6]This definition is inspired by the definition of *depth* in [17, 72] where the depth of the segment $p_a p_b$ is the minimum number of points of $P \setminus \{p_a, p_b\}$ on either side of the line $p_a p_b$.

Theorem 3.2.7 ([72]). *Any polygon S with n segments and endpoints P in convex position admits (in the* `Convex Cycle` *version) an untangle sequence of length at most $2n - 7 = O(n)$ for $n \geq 4$.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Cycle}}(n) \leq 2n - 7 = O(n) \textit{ for } n \geq 4.$$

The proof of Theorem 3.2.7 relies on the following lemma.

Lemma 3.2.8 ([72]). *Any polygon S with n segments and endpoints P in convex position admits an uncrossable segment $s \in S$ after at most 1 flip (in the* `Convex Cycle` *version).*

Proof. Note that, in the case of polygons with endpoints in convex position, the uncrossable segments are the segments contained in the boundary of the convex hull of P which are the segments of null crossing depth. We assume that S contains no uncrossable segment. Let $p_1, \ldots, p_{|P|}$ be the points of P sorted in counterclockwise order along the convex hull boundary.

Let $p_a p_b \in S$ $(a < b)$ be a segment *with crossings* (hence, of crossing depth at least one) minimizing $\delta_\times(p_a p_b)$. Let $q_1, \ldots, q_{\delta_\times(p_a p_b)}$ be the points defining $\delta_\times(p_a p_b)$ sorted in counterclockwise order along the convex hull boundary.

Since $p_a p_b$ has minimum positive crossing depth, and since we assume that S contains no uncrossable segment, the point q_1 is the endpoint of two segments s, s', each of them either crossing $p_a p_b$. It is easy to see that, by definition of a flip in the `Cycle` version, flipping one of these two segments, say s, with the segment $p_a p_b$ inserts the segment $q_1 p_a$ while flipping s' with the segment $p_a p_b$ does not insert the segment $q_1 p_a$. We choose to remove the segments s and $p_a p_b$, because, being of null crossing depth, the inserted segment $q_1 p_a$ is uncrossable. □

Proof of Theorem 3.2.7. We now prove Theorem 3.2.7 by induction on n.

If $n = 4$, then $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Cycle}}(n) = 1$.

If $n \geq 5$, and if $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Cycle}}(n-1) \leq 2(n-1) - 7$, then, by Lemma 3.2.8, using at most 1 flip ensures that there exists at least one uncrossable segment s in S. We invoke the induction hypothesis on the polygon with $n - 1$ segments obtained from S by contracting the segment s to a point. The resulting flip sequence also transforms S into a polygon where only the two segments adjacent to s may cross. A final flip is enough to fully untangle S. The total number of flips use to untangle S is at most $2 + 2(n-1) - 7 = 2n - 7$, concluding the induction. □

The following theorem refines Theorem 3.2.7 with a tight bound.

Theorem 3.2.9 ([85]). *Any polygon with $n \geq 7$ segments and endpoints in convex position admits (in the* `Convex Cycle` *version) an untangle sequence of length at most $n - 2 = O(n)$.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Cycle}}(n) \leq n - 2 = O(n) \quad \textit{for } n \geq 7.$$

The article [85] is entirely dedicated to the proof of Theorem 3.2.9; we omit this proof here.

The following theorem also provides a linear upper bound, but in the `Convex Bipartite Matching` version.[7]

Theorem 3.2.10 ([17]). *Consider a set S of n segments with endpoints P in convex position partitioned into n red points and n blue points such that (P, S) forms a bipartite matching.*

In the `Convex Bipartite Matching` *version, there exists an untangle sequence of S of length at most $2n - 3 = O(n)$.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\mathrm{R}}_{\texttt{Convex Bipartite Matching}}(n) \leq 2n - 3 = O(n).$$

To prove Theorem 3.2.10, we use removal choice to iteratively remove a segment of minimal positive crossing depth and insert a segment of null collapsed crossing depth using two flips.

Proof of Theorem 3.2.10. Consider a multiset S of n segments with endpoints P in convex position. Let $p_1, \ldots, p_{|P|}$ be the points of P sorted in counterclockwise order along the convex hull boundary.

We repeat the following procedure until there are no more crossings. Let $p_a p_b \in S$ $(a < b)$ be a segment *with crossings* (hence, of crossing depth at least one) minimizing $\delta_\times(p_a p_b)$. Without loss of generality, assume that p_a is red and that p_b is blue. Let $q_1, \ldots, q_{\delta_\times(p_a p_b)}$ be the points defining $\delta_\times(p_a p_b)$ sorted in counterclockwise order along the convex hull boundary.

We have the following two cases.

Case 1: using one flip. In Case 1, q_1 is blue or $q_{\delta_\times(p_a p_b)}$ is red. We assume without loss of generality that q_1 is red. Since $p_a p_b$ has minimum positive crossing depth, the point q_1 is the endpoint of a segment $q_1 p_c \in S$ that crosses $p_a p_b$. In this case, we chose to remove the segments $q_1 p_c$ and $p_a p_b$. By definition of q_1, the crossing depth of the inserted segment $q_1 p_a$ is 0.

Case 2: using two flips. In Case 2, q_1 is red and $q_{\delta_\times(p_a p_b)}$ is blue. Let i be the smallest index such that q_i is blue. Thus, q_{i-1} is red. Again, since $p_a p_b$ has minimum positive crossing depth, the point q_{i-1} is the endpoint of a segment $q_{i-1} p_c \in S$ that crosses $p_a p_b$. We chose to remove the segments $q_{i-1} p_c$ and $p_a p_b$. The inserted segments are $q_{i-1} p_b$ and $p_a p_c$.

At this point, the segment $q_{i-1} p_b$ crosses a segment $q_i p_d$ incident to q_i (in fact, any such segment). We chose to remove these two segments, i.e., $q_{i-1} p_b$ and $q_i p_d$. The inserted segments are $q_{i-1} p_b$ and $p_a p_c$. The crossing depth of the inserted segment $q_i q_{i-1}$ is 0, by definition of i.

In both cases, we have used at most 2 flips to remove segments with positive crossing depths and insert at least one segment with null crossing depth. Thus, the number of segments with positive crossing depth decreases by at least 1 every two flips. Using

[7] The upper bound proved in [17] is $2n - 2$.

at most $2(n-2)$ flips is enough to reach $n-2$ segments of crossing depth 0 (which no longer participate in a flip by Lemma 2.6.1 [19, 31]). The last two segments with positive crossing depths are crossing, and flipping them necessarily inserts two segments of crossing depth 0. The claimed bound follows. $\square$

The following theorem provides a quasilinear upper bound in the `Convex Tree` version.

Theorem 3.2.11 ([17]). *Consider a set S of n segments with endpoints P in convex position such that (P, S) forms a tree.*

In the `Convex Tree` *version, there exists an untangle sequence of S of length $O(n \log n)$.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Tree}}(n) = O(n \log n).$$

As we prove the more general Theorem 5.2.1 in the contribution part of this dissertation using some simpler arguments than the proof in [19], we do not provide a proof of Theorem 3.2.11 here.

3.2.3 Untangling with Both Choices

In this section, we prove a linear lower bound on $\mathbf{d}^{\mathtt{RI}}_{\mathtt{Convex\ Matching}}$ and a linear upper bound on $\mathbf{d}^{\mathtt{RI}}_{\mathtt{Convex\ Multigraph}}$ from the literature.

Lower Bound

In this section, we prove the following lower bound.

Theorem 3.2.12 ([19]). *There exists a set S of n segments with endpoints P in convex position such that (P, S) forms a matching (respectively a bipartite matching, a cycle, a tree), and such that (in the* `Convex Matching` *version, respectively in the* `Convex Bipartite Matching` *version, in the* `Convex Cycle` *version, in the* `Convex Tree` *version) any untangle sequence of S is of length at least $n - 1 = \Omega(n)$ (respectively $n - 1 = \Omega(n)$, $\frac{n}{2} - 1 = \Omega(n)$ for even n, $\frac{n-1}{2} = \Omega(n)$ for odd n).*

In other words, we have the following lower bounds:

$$\begin{aligned}
\mathbf{d}^{\mathtt{RI}}_{\mathtt{Convex\ Matching}}(n) &\geq n - 1 = \Omega(n), \\
\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Bipartite\ Matching}}(n) &\geq n - 1 = \Omega(n), \\
\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Cycle}}(n) &\geq \frac{n}{2} - 1 = \Omega(n) \text{ for even } n, \\
\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Tree}}(n) &\geq \frac{n-1}{2} = \Omega(n) \text{ for odd } n.
\end{aligned}$$

In the following, first we define a class of sets of segments called *combs*, then we prove Theorem 3.2.12 by proving that any untangle sequence of an n-comb is of length exactly $n - 1$ using splitting to invoke induction.

Combs. We say that a set S of n segments with endpoints P in convex position such that (P, S) forms a bipartite matching is an *n-comb* if there exists a segment $s \in S$ such that s crosses all the segments in $S \setminus \{s\}$ and such that $S \setminus \{s\}$ is crossing free.

Proof of Theorem 3.2.12. Let S be an n-comb. We prove by induction on n that the length $f(n)$ of any untangle sequence of S is exactly $n - 1$.

If $n = 1$, then S contains only s and the claim $f(1) = 0$ is trivial.

If $n \geq 2$, then any flip removes s and some other segment $s' \in S \setminus \{s\}$. Independently of which pair of segments is inserted, the line ℓ containing s' does not cross any segment of S after the flip. Thus, ℓ splits S into an n_1-comb S_1 and an n_2-comb S_2, for some $n_1 \geq 1$ and $n_2 \geq 1$ such that $n_1 + n_2 = n$. This yields the following recurrence relation

$$f(n) = 1 + f(n_1) + f(n_2),$$

which easily solves using induction hypothesis to

$$f(n) = n - 1,$$

concluding the induction.

This exact same upper bound holds on to $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Bipartite\ Matching}}(n)$ by the transfer reduction (Lemma 2.3.3). It is easy to transform an n-comb into a cycle by adding n segments on the convex hull boundary of P. These new segments are uncrossable (Lemma 2.6.1) thus do not interfere. The insertion choice is now set to preserve the `Cycle` property, but the induction above makes no assumption about insertion choice. Therefore, the proof holds in the `Convex Cycle` version too. Similarly, it easy to remove one of these extra segments to get a path (which is a tree), and to extend the same proof to the `Convex Tree` version. □

Upper Bound

In this section, we prove the following upper bound. It is originally proven in [17] for the `Convex Matching` version.

Theorem 3.2.13 ([17])**.** *Consider a multiset S of n segments with endpoints P in convex position.*

In the `Convex Multigraph` *version, there exists an untangle sequence of S of length at most $n - 1 = O(n)$.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\mathtt{RI}}_{\mathtt{Convex\ Multigraph}}(n) \leq n - 1 = O(n).$$

In the following, we prove Theorem 3.2.13 using removal and insertion choices to iteratively remove a segment of minimal positive crossing depth (defined before Theorem 3.2.7) and insert a segment of null collapsed crossing depth using one flip.

Proof of Theorem 3.2.13. Consider a multiset S of n segments with endpoints P in convex position. Let $p_1, \ldots, p_{|P|}$ be the points of P sorted in counterclockwise order along the convex hull boundary.

We repeat the following procedure until there are no more crossings. Let $p_a p_b \in S$ $(a < b)$ be a segment *with crossings* (hence, of crossing depth at least one) minimizing

$\delta_\times(p_a p_b)$. Let $q_1, \ldots, q_{\delta_\times(p_a p_b)}$ be the points defining $\delta_\times(p_a p_b)$ sorted in counterclockwise order along the convex hull boundary.

Since $p_a p_b$ has minimum positive crossing depth, the point q_1 is the endpoint of a segment $q_1 p_c \in S$ that crosses $p_a p_b$. We chose to remove the segments $q_1 p_c$ and $p_a p_b$, and to insert the segments $q_1 p_a$ and $p_b p_c$.

By definition of q_1, the crossing depth of the inserted segment $q_1 p_a$ is 0. We have thereby removed two segments with positive crossing depths and inserted at least one segment with null crossing depth. Thus, the number of segments with positive crossing depth decreases by at least 1 at each flip.

Using at most $n-2$ flips is enough to reach $n-2$ segments of crossing depth 0 (which no longer participate in a flip by Lemma 2.6.1 [19, 31]). The last two segments with positive crossing depths are crossing, and flipping them necessarily inserts two segments of crossing depth 0. The claimed bound follows. □

3.3 Red-on-a-Line Point Sets

In this section, we present bounds on the length of untangle sequences in the `Redonaline Matching` version, i.e., in the `Bipartite Matching` version in the special case where the bipartite point set has all the points of one of its partitions, say the red points, on a common line defining the x-axis.

The x-axis splits any bipartite matching on such a point set. The splitting lemma (Lemma 2.6.1 [19, 31]) allows to study flips above the x-axis and flips below separately. We thus assume that all the blue points lie above the x-axis. We call such a point set a *red-on-a-line* point set.

This special type of point set may be viewed as a transition between the well understood convex bipartite point sets where the red points are on one line ℓ_1 and the blue points are on an other line ℓ_2 parallel to ℓ_1 (this case is fully encoded by permutations and inversions; see Theorem 3.2.1) and the general bipartite point sets.

However, only the upper bound on the length of untangle sequences with removal choice is specific to general red-on-a-line point sets.

3.3.1 Untangling with Removal Choice

In this section, we prove a quadratic upper bound on $\mathbf{d}^{\mathtt{R}}_{\texttt{Redonaline Matching}}$ from the literature.

Upper Bound

In this section, we prove the following upper bound.

Theorem 3.3.1 ([17]). *Consider a set S of n segments with endpoints P partitioned into n red points and n blue points such that the red points lie on the x-axis, such that the blue points lie above the x-axis, and such that (P, S) forms a bipartite matching.*

In the `Redonaline Matching` *version, there exists an untangle sequence of S of length at most $n(n-1) = O(n^2)$.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\mathtt{R}}_{\texttt{Redonaline Matching}}(n) \leq n(n-1) = O(n^2).$$

The proof of Theorem 3.3.1 invoke the following lemma as a subroutine.

Lemma 3.3.2. *Consider a set S of n segments with endpoints P partitioned into n red points and n blue points such that the red points lie on the x-axis, such that all the blue points lie above the x-axis, and such that (P, S) forms a bipartite matching.*

If the highest blue point b is matched to the left most red point r in S, and if all the crossings of S are on the segment rb, then (in the `Redonaline Matching` *version) there exists an untangle sequence of S of length at most $n - 1 = O(n)$.*

Proof. We prove Lemma 3.3.2 by induction on n. If $n = 1$, then Lemma 3.3.2 holds.

If $n > 1$, then we show that one flip is enough to define a line partitioning the segments into two subsets S_1 and S_1' such that S_1' is not empty and has no crossing.

Let b' be the first blue point in counterclockwise order around the point r (the ray from r including in the x-axis going right is the origin of the angles). Let r' be the red point matched to b' in S. We perform the flip removing the segments $rb, r'b'$ and inserting $rb', r'b$.

After this flip (modifying S in place), the line $r'b'$ does not cross any segment of S therefore partitioning the segments into two subsets. Let S_1 be the segments of S lying on the left of the line $r'b'$ and let S_1' be the segments of S lying on the right of the line $r'b'$. The set of segments S_2 is not empty since it contains at least the segment $r'b$. By construction of the point b', the segment $r'b$ is crossing free. Thus, the segments in S_2 are all crossing free.

By induction hypothesis, we untangle S_1 using $n - 2$ flips, concluding the proof of Lemma 3.3.2. □

Note that the bound of Lemma 3.3.2 is tight. Indeed, there exists a red-on-a-line matching with n segments such that any removal strategy leads to the same untangle sequence of length $n - 1$. In fact, any removal strategy is futile as there is exactly one crossing pair of segments before each flip of the only possible untangle sequence (see Figure 3.5 [19]). An alternative proof of this tightness (also holding for convex position) is given by Theorem 3.2.12.

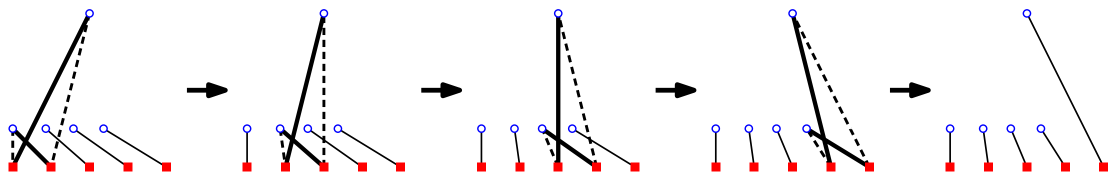

Figure 3.5: A lower-bounding untangle sequence of a red-on-a-line matching where only one pair of segments is crossing at each step, thus imposing removal choice. This figure appears in [19].

We are now ready to prove Theorem 3.3.1.

Proof of Theorem 3.3.1. We describe an iterative algorithm to untangle S. We label the red points from left to right as follows: $r_1, \ldots, r_n$.

For i from 1 to n, we perform Phase 1 then Phase 2 described next.

Loop invariant. Entering the i-th loop, all the segments incident to the red points on the left of r_i are crossing free.

Phase 1: find a blue point b and insert the segment $r_i b$ (using at most $n-i$ flips) such that b is the highest blue point, r is the leftmost red point of a subset S' of S. We define b as the blue point we reach at the end of the following path u. We start the path at r_i and we follow the segment incident to r_i until we encounter a segment of S. We turn left and continue the path on this new segment. We repeat, systematically turning left at every crossing point until we reach a blue point, which we name b.

Let s and s' be the last two segments we visited in the path u. Repeatedly flipping s, s' insert the segment $r_i b$ after at most $n-i$ flips.

We now define S'. Let r_0 and b_0 be dummy red and blue points such that r_0 lies on the x-axis on the left of all the other red points, such that b_0 becomes the highest blue point, and such that the segment $r_0 b_0$ is crossing-free. We shoot a horizontal ray from b heading leftwards. Let r_j be the red endpoint of the first segment of $S \cup \{r_0 b_0\}$ hit by this ray, and let q be the hitting point. We finally define S' the segments of S which are in the interior of the quadrilateral $q r_j r_i b$.

Phase 2: invoke Lemma 3.3.2 on S' (using at most $i-1$ flips). The quadrilateral $q r_j r_i b$ is convex and is not intersected by any segment of S. Therefore, after having applied Lemma 3.3.2, the segments of S' are not only not crossing each other but also not crossing other segments of $S \setminus S'$. This ensures that the loop invariant holds.

To conclude the proof of Theorem 3.3.1, an upper bound to the total number of flips used by the algorithm is given by the following sum.

$$\sum_{i=1}^{n} (n-i) + (i-1) = n(n-1) \qquad \square$$

Chapter 4

Untangling with No Choice

In this chapter, we study $\mathbf{d}^{\emptyset}$, the length of the longest untangle sequences in several versions. First, we prove reductions between several versions of $\mathbf{d}^{\emptyset}$. Then, we prove a lower bound, an upper bound parameterized by t (the sum of the degrees of the points which are not in convex position with the rest), and an upper bound in the `Redonaline` version. 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.

4.1 Reductions

In this section, we provide a series of inequalities relating the different versions of $\mathbf{d}^{\emptyset}$. In particular, we show that all the different versions of $\mathbf{d}^{\emptyset}$ have the same asymptotic behavior.

Theorem 4.1.1 ([30]). *For all positive integer n and for* `Π` *being either the empty property or the* `Convex` *property, we have the following relations:*

$$\mathbf{d}^{\emptyset}_{\Pi\ \texttt{Matching}}(n) = \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Multigraph}}(n), \tag{4.1}$$

$$2\mathbf{d}^{\emptyset}_{\Pi\ \texttt{Matching}}(n) \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Bipartite Matching}}(2n) \ \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Matching}}(2n), \tag{4.2}$$

$$2\mathbf{d}^{\emptyset}_{\Pi\ \texttt{Bipartite Matching}}(n) \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Cycle}}(3n) \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Matching}}(3n), \tag{4.3}$$

$$2\mathbf{d}^{\emptyset}_{\Pi\ \texttt{Bipartite Matching}}(n) \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Tree}}(3n) \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Matching}}(3n). \tag{4.4}$$

Proof. Equality (4.1) can be rewritten $\mathbf{d}^{\emptyset}_{\Pi\ \texttt{Multigraph}}(n) \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Matching}}(n) \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Multigraph}}(n)$. Hence, we have to prove eight inequalities. The right-side inequalities follow from Lemma 2.3.2, since the left-side property is stronger than the right-side property (using the equality (4.1) for inequalities (4.3) and (4.4)).

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.

We first prove all the inequalities assuming that `Π` is the empty property. Preserving the `Convex` property follows from the fact that the extra copy a point added by the following reductions is placed in convex position while still being close enough to its original location (Figure 4.1(a)).

Proving $\mathbf{d}^{\emptyset}_{\Pi\ \texttt{Multigraph}}(n) \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Matching}}(n)$ (4.1). We prove the left inequality of (4.1). A point of degree d° larger than 1 can be replicated as d° 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 4.1(a)). Thus, for any flip sequence in the `Multigraph` version, there exists a flip sequence in the `Matching` version of equal length, yielding $\mathbf{d}^{\emptyset}_{\Pi\ \texttt{Multigraph}}(n) \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Matching}}(n)$.

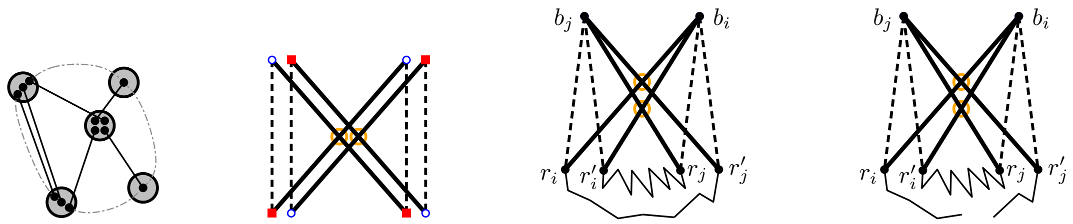

(a) (b) (c) (d)

Figure 4.1: (a) A crossing-free multigraph transformed into a matching by replacing multi-degree points by clusters of points. This transformation (like the next ones) preserves the property that a set of points is in convex position (their convex hull boundary is dashed-dotted). (b) Two flips in the `Bipartite Matching` version simulating one flip in the `Matching` version. (c) Two flips in the `Cycle` version simulating one flip in the `Bipartite Matching` version. (d) Two flips in the `Tree` version simulating one flip in the `Bipartite Matching` version.

Proving $2\mathbf{d}^{\emptyset}_{\Pi\ \texttt{Matching}}(n) \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Bipartite Matching}}(2n)$ (4.2). The left inequality of (4.2) 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 bipartite 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 4.1(b) shows this construction. A sequence of k flips on S provides a sequence of $2k$ flips on S'. Hence, $2\mathbf{d}^{\emptyset}_{\Pi\ \texttt{Matching}}(n) \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Bipartite Matching}}(2n)$.

Proving $2\mathbf{d}^{\emptyset}_{\Pi\ \texttt{Bipartite Matching}}(n) \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Cycle}}(3n)$ (4.3). To prove the left inequality of (4.3), we start from a bipartite 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 $r_1, b_1, r'_1, \ldots, r_i, b_i, r'_i, \ldots, r_n, b_n, r'_n \ldots$ where r_i is matched to b_i in S (Figure 4.1(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_ib_i, r_jb_j and inserting the pair of segments r_ib_j, r_jb_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 $r_1, b_1, r'_1, \ldots, r_i, b_j, r'_i, \ldots, r_j, b_i, r'_j, \ldots, r_n, b_n, r'_n, \ldots$ on which we can apply the next flips in the same way. Hence, $2\mathbf{d}^{\emptyset}_{\Pi\ \texttt{Matching}}(n) \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Cycle}}(3n)$.

Proving $2\mathbf{d}^{\emptyset}_{\Pi\ \texttt{Bipartite Matching}}(n) \leq \mathbf{d}^{\emptyset}_{\Pi\ \texttt{Tree}}(3n)$ (4.4). The proof of the left inequality of (4.4) follows the exact same construction as in the proof of the left inequality of (4.3). 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 4.1(d)). □

4.2 Lower Bound for Red-on-a-Line Matchings

In this section, we prove the following quadratic lower bound. Although this lower bound only improves the state of the art by a constant factor, its proof nonetheless exhibits the longest untangle sequence ever know. This untangle sequence is 1.5 times longer than the longest untangle sequence in the `Convex Bipartite Multigraph` version. To achieve this 1.5 factor, we take advantage of the non-convex aspect in a non-intuitive fashion.

Theorem 4.2.1 ([33])**.** *In the* `Redonaline Matching` *version for even n, there exists a set S of n segments with endpoints P partitioned into n red points and n blue points such that the red points lie on the x-axis, such that the blue points lie above the x-axis, such that (P,S) forms a bipartite matching, and such that there exists an untangle sequence of S of length $\frac{3}{2}\binom{n}{2} - \frac{n}{4} = \Omega(n^2)$.*

In other words, for even n, we have the following lower bound:

$$\mathbf{d}^{\emptyset}_{\texttt{Redonaline Matching}}(n) \geq \frac{3}{2}\binom{n}{2} - \frac{n}{4} = \Omega(n^2).$$

To prove Theorem 4.2.1, it suffices to present a long untangle sequence. The initial matching of the sequence is represented in Figure 4.2(a).

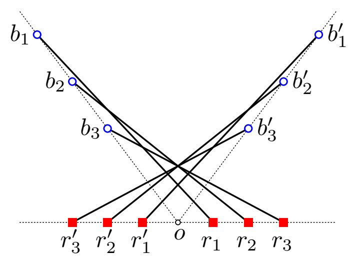

Figure 4.2: A 3-butterfly to lower bound $\mathbf{d}^{\emptyset}(6)$.

We provide a $2m$-segment red-on-a-line matching which we call an *m-butterfly*. There exists an untangle sequence starting at an m-butterfly of length $\frac{3}{2}\binom{2m}{2} - \frac{m}{2}$. Next, we give the precise definition of an m-butterfly and some of its important properties. Then, we give some intuition of how to come up with an untangle sequence longer than the number of pairs of segments. Finally, we prove that there exists an untangle sequence starting at an m-butterfly of length $\frac{3}{2}\binom{2m}{2} - \frac{m}{2}$ with two lemmas.

Butterfly. For an integer m, we define an *m-butterfly* as the following matching with $n = 2m$ segments. For i from 1 to m we have red points $r_i = (i/(m+1), 0)$ and $r'_i = (-i/(m+1), 0)$ as well as blue points $b_i = (i-(m+1), (m+1)-i)$ and $b'_i = ((m+1)-i, (m+1)-i)$. We match r_i to b_i and r'_i to b'_i. Next, we discuss important properties of an m-butterfly.

We call a red-on-a-line *convex* matching an *n-star* if all the $\binom{n}{2}$ pairs of segments cross. We say that an n-star *looks* at a point p if the blue points are all on a common line, and if p is the intersection of this line with the line of the red points. We also say that two red-blue point sets R, B and R', B' are *fully crossing* if all the pairs of segments of the form $\{rb, r'b'\}$ cross, where $(r, b, r', b') \in R \times B \times R' \times B'$. Two matchings are fully crossing if their underlying red-blue point sets are fully crossing. An m-butterfly is a red-on-a-line matching consisting of two fully crossing m-stars both looking at the same point $o = (0, 0)$ (Figure 4.2(a) represents these properties but it is not drawn to scale).

Intuition. In the following, we use the state tracking framework from Section 5.8 to describe how to come up with an untangle sequence starting at an m-butterfly with more than $\binom{2m}{2}$ flips. We consider a sequence of tracking choices with no $\mathbf{H} \to \mathbf{X}$ transition (Lemma 5.8.2) for the long untangle sequence we build. We take advantage of the non-convex position of the blue points to create flip situations such as in Figure 5.11(a), where an **H**-pair is turned into a **T**-pair.

For instance, let us consider an **X**-pair of one of the m-stars composing the m-butterfly. At some point of the untangle sequence, we flip this **X**-pair, turning it into an **H**-pair. Later on, we turn this **H**-pair into a **T**-pair, as in Figure 5.11(a). Still later on, we turn this **T**-pair into an **X**-pair again, similarly to the pairs involving the smallest horizontal segment in Figure 1.2. This **X**-pair will be flipped again.

We manage to carry out this whole process to flip twice all the $2\binom{m}{2}$ pairs of the two m-stars composing the m-butterfly while still having one flip for every other pair. In total, we reach $\frac{3}{2}\binom{2m}{2} - \frac{m}{2}$ flips.

Proof of Theorem 4.2.1. We exhibit an untangle sequence starting at any m-butterfly of length $\frac{3}{2}\binom{2m}{2} - \frac{m}{2}$ (Figure 4.3 and Figure 4.4). The untangle sequence can be divided into two phases.

The first phase consists of (i) $\binom{m}{2}$ flips applied to the m-star submatching defined by the m leftmost red points (see the proof of Theorem 3.2.1 [19], and Figure 4.3, steps 0 to 3), and of (ii) $\binom{m}{2}$ more flips applied to the m rightmost red points (Figure 4.3, steps 3 to 6). At this point, we have two sets of m crossing-free segments, each set fully crossing the other.

The second phase repeats m times the following routine.

1. Flip the segments defined by the innermost red points r_1 and r'_1 (Figure 4.3 and 4.4, steps 6 to 7, 11 to 12, and 16 to 17). After this flip, the submatching defined by the m leftmost red points and their matched points consists of $m-1$ crossing-free segments intersected by the segment from r'_1. A similar statement holds for the submatching defined by the m rightmost red points.

2. Untangle the submatching defined by the m leftmost red points with $m-1$ flips in the following manner (Figure 4.3 and 4.4, steps 7 to 9, 12 to 14, and 17 to

19). Flip of the two crossing segments with the rightmost red points, say r'_k and r'_{k+1} with $k \in \{1, \ldots, m-1\}$, and repeat. Such a flip produces a segment from r'_{k+1} crossing the segments whose red points are on the left of r_{k+1}, and an other segment from r'_k crossing none of segments of the submatching. The number of crossings in the submatching decreases by 1 at each flip.

3. Similarly, untangle the m rightmost red points with $m-1$ more flips (Figure 4.3 and 4.4, steps 9 to 11, 14 to 16, and 19 to 21).

Each loop decreases the number of "long" segments (i.e., segments joining one of the leftmost red points to one of the rightmost blue points, or vice-versa) by 2. At the end of the process, the left submatching is crossing-free; so is the right one; and the two of them do not intersect anymore.

Summing up, the total length of the untangle sequence is $2\binom{m}{2} + 2m(m-1) + m$. Simple calculation yields the lemma. □

4.3 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 translations to the other versions follows from the reductions in Theorem 4.1.1.

Theorem 4.3.1 ([30]). *Consider a multiset S of n segments with endpoints P partitioned into $P = C \cup T$ where C is in convex position. We define the parameter t as the sum of the degrees of the points in T.*

In the `Multigraph` version, any untangle sequence of S is of length $O(tn^2)$.

In other words, we have the following upper bound:

$$\mathbf{d}^{\emptyset}_{\texttt{Multigraph}}(n,t) = O(tn^2).$$

Proof. The proof strategy is to combine the potential χ used in the proof of Theorem 3.2.2 [17] with the potential Λ_L used in the proof of Theorem 3.1.3 [82]. 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. $\chi(f(S)) \geq \chi(S)$ (as shown in Figure 1.2).

The potential Λ_L is derived from the line potential introduced in [83] 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.

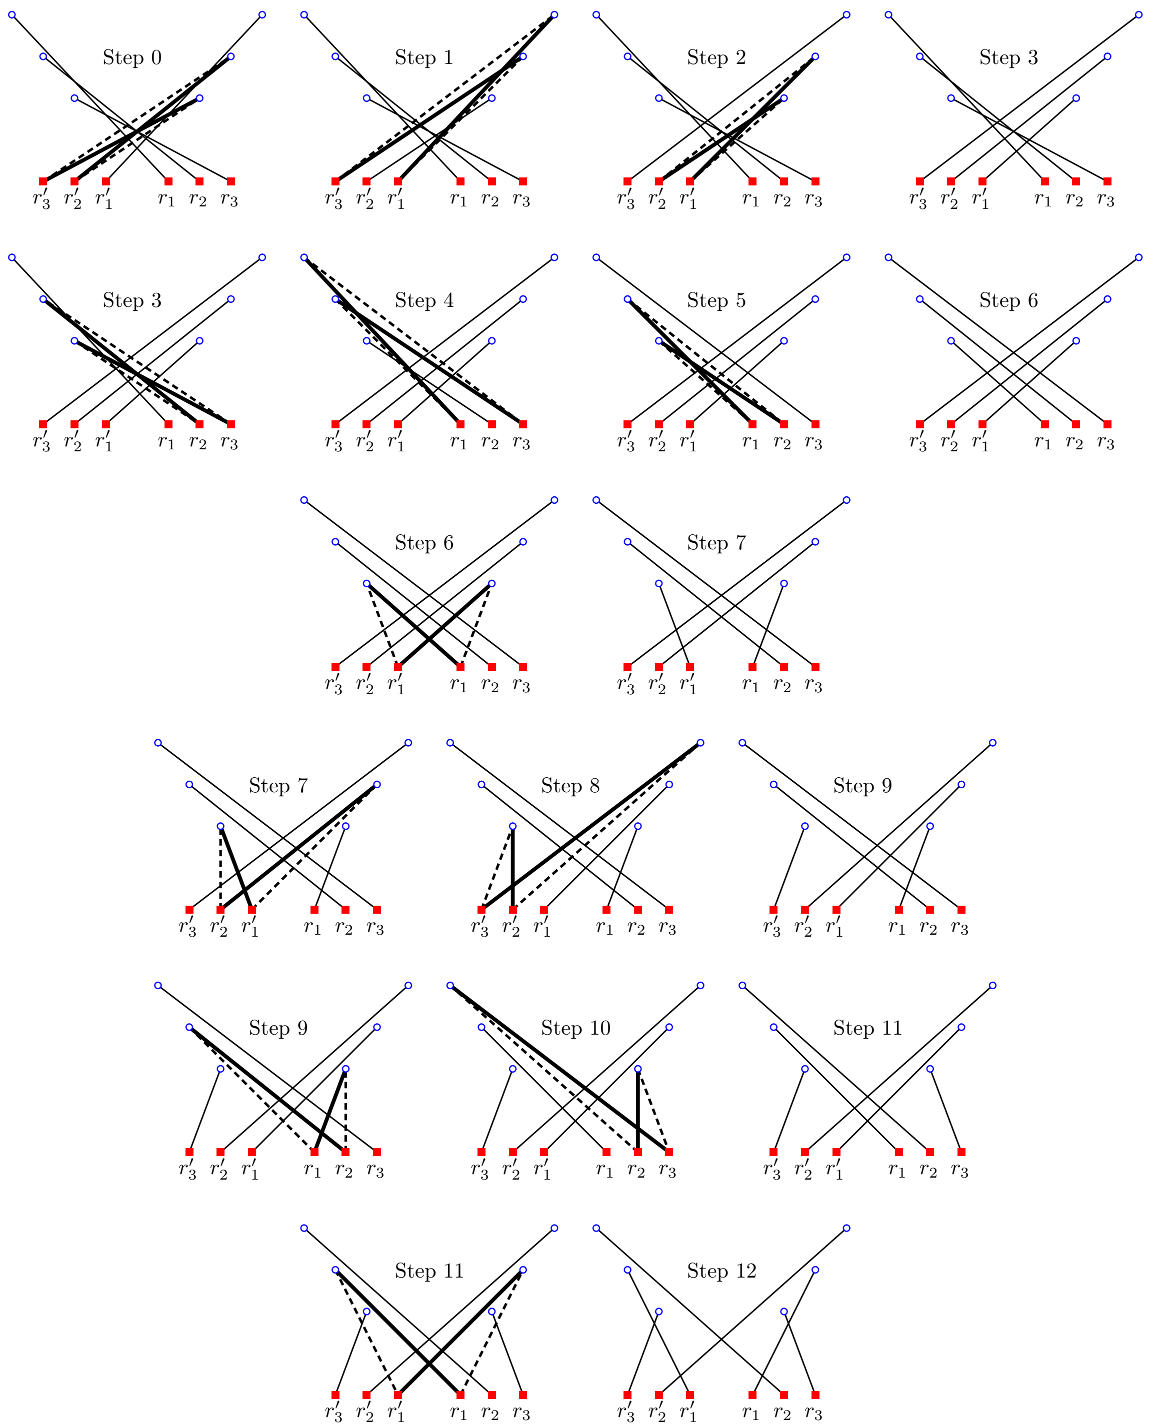

Figure 4.3: First part of an untangle sequence starting at a 3-butterfly of length 21 illustrating Theorem 4.2.1 and its proof. Each line corresponds to a portion of the proof, with repetitions added for clarity.

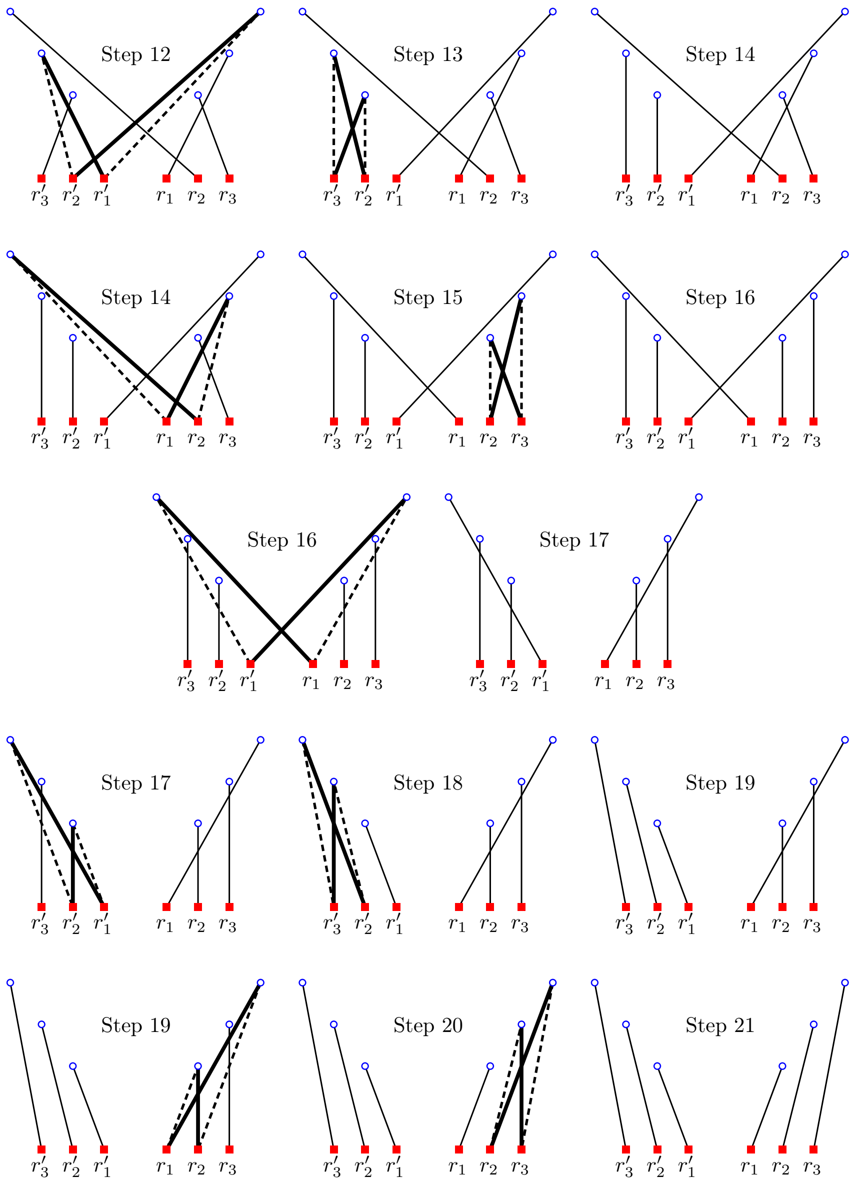

Figure 4.4: Second part of an untangle sequence starting at a 3-butterfly of length 21 illustrating Theorem 4.2.1 and its proof. Each line corresponds to a portion of the proof, with repetitions added for clarity.

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: $0 \le \Phi(S) \le O(tn^2)$. To complete the proof of Theorem 4.3.1, we show that any flip decreases Φ by at least 1.

We consider an arbitrary flip f removing the pair of segments p_1p_3, p_2p_4 and inserting the pair of segments p_1p_4, p_2p_3. Let o be the point of intersection of p_1p_3 and p_2p_4. It is shown in Lemma 3.1.1 [83] 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_1p_4 and p_2p_3 and exactly one of the four flipped points belongs to ℓ. We call these lines *f-critical* (Figure 4.5(a)).
- The potential Λ_ℓ decreases by 2 if the line ℓ strictly separates the final segments p_1p_4 and p_2p_3. We call these lines *f-dropping* (Figure 4.5(b)).
- The potential Λ_ℓ remains stable in the remaining cases.

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 4.5(a)).

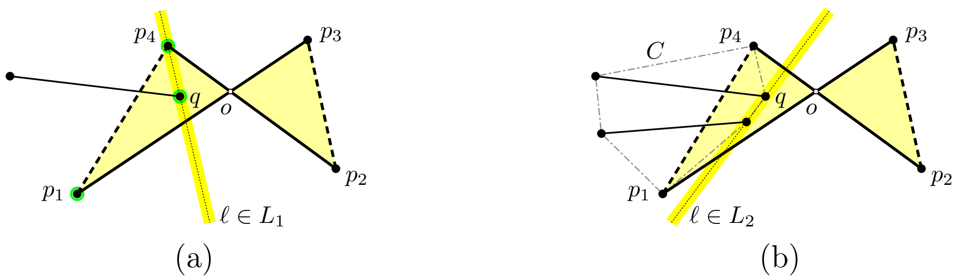

Figure 4.5: (a) An f-critical line ℓ for a flip f removing p_1p_3, p_2p_4 and inserting p_1p_4, p_2p_3. This situation corresponds to case (2a) with $\ell \in L_1$. (b) An f-dropping line ℓ. This situation corresponds to case (2b) with $\ell \in L_2$.

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 $0 \le k \le n-1$, 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 $(0 \le i \le k)$, inside the non-simple polygon p_1, p_4, p_2, p_3 (Figure 4.5). Next, we show that each point $q \in \{q_0, \ldots, q_k\}$ 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 4.5(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 4.5(b)) or f-critical (the line qp_4 in

Figure 4.5(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 $p_1 o p_4$, and the theorem follows. □

4.4 Upper Bound for Red-on-a-Line Matchings

In this section we prove the following theorem.

Theorem 4.4.1 ([33]). *Consider a set S of n segments with endpoints P partitioned into n red points and n blue points such that the red points lie on the x-axis, such that the blue points lie above the x-axis, and such that (P, S) forms a bipartite matching.*

In the `Redonaline Matching` *version, any untangle sequence of S is of length at most $\binom{n}{2}\frac{n+4}{6} = O(n^3)$.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\emptyset}_{\texttt{Redonaline Matching}}(n) \leq \binom{n}{2}\frac{n+4}{6} = O(n^3).$$

To prove Theorem 4.4.1, we define a potential function Φ that maps a red-on-a-line matching to an integer from 0 to $\binom{n}{2}\frac{n+4}{3}$. Since Φ decreases by at least 2 at each flip, the theorem follows. We first give the definitions needed to present Φ. Then, we prove four lemmas yielding Theorem 4.4.1.

Let S be a red-on-a-line matching. Let $r_1, \ldots, r_n$ be the red points, from left to right. Let ℓ be a line, parallel to the line of the red points and above all the points. For each k in $\{1, \ldots, n\}$, we project the blue points onto ℓ, using r_k as a focal point. More precisely, each blue point b maps to a point $t_k(b)$, the intersection between the ray $r_k b$ and the line ℓ (Figure 4.6(a)). We also define the function t_k of a red-blue segment rb as the segment $t_k(rb) = rt_k(b)$ (Figure 4.6(b)).

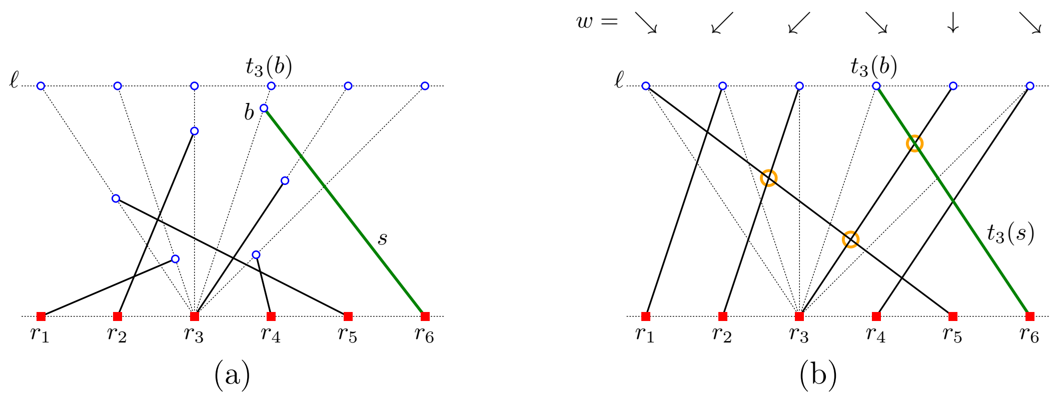

Figure 4.6: (a) The projection t_k for $k = 3$. (b) The segments $t_3(\cdot)$. The three 3-observed crossing 3-pairs are circled.

We may abbreviate a pair of segments $r_i b, r_j b'$ as $\langle i, j \rangle$ when the points b and b' can be deduced from the underlying matching. Let k be an integer in $\{1, \ldots, n\}$. We say that two segments are *k-observed crossing* if the extended projection $t_k(\cdot)$ maps them to crossing segments (Figure 4.6(b)). A pair of segments $\langle i, j \rangle$ is a *k-pair* if $i \leq k \leq j$. A *k-flip* is then a flip of a k-pair. We have the following lemma.

Lemma 4.4.2. *A crossing k-pair is necessarily k-observed crossing.*

Proof. Let $r_i b, r_j b'$ be a crossing k-pair. We suppose, without loss of generality, that $i < j$ (e.g. $i = 2$, $k = 3$, and $j = 5$ in Figure 4.6).

The fact that the k-pair $r_i b, r_j b'$ is crossing means that the four points are in convex position, and that they appear as r_i, r_j, b, b' on their convex hull in counter-clockwise order. Since $i \le k \le j$, the point r_k is also on the boundary of the convex hull of the four points. Therefore, the projection $t_k(\cdot)$ will not change the convex-hull order and the segments $r_i t_k(b)$ and $r_j t_k(b')$ will cross. $\square$

We define $\Phi_k(S)$, the *k-th potential* of S, as the number of k-observed crossing k-pairs (Figure 4.6(b)). Lemma 4.4.3 shows that the k-th potential Φ_k is at most $(k-1)(n-k)+n-1$. Lemma 4.4.4 shows that Φ_k never increases, and decreases by at least 1 at each k-flip.

Lemma 4.4.3. *The k-th potential Φ_k takes integer values from 0 to $k(n+1)-k^2-1$.*

Proof. The k-th potential $\Phi_k(S)$ is at most the number of k-pairs in S, crossing or not. There are exactly $(k-1)(n-k)$ k-pairs of the form $\langle i, j\rangle$ with $i < k < j$. There are exactly $k-1$ k-pairs of the form $\langle i, k\rangle$ with $i < k$. There are exactly $n-k$ k-pairs of the form $\langle k, j\rangle$ with $k < j$. In total, there are $k(n+1)-k^2-1$ k-pairs in S. $\square$

Lemma 4.4.4. *The k-th potential Φ_k never increases, and decreases by at least 1 at each k-flip.*

Proof. We order the projected blue points on ℓ from left to right. We then map each projected blue point $t_k(b)$ to an element in $\{\swarrow, \downarrow, \searrow\}$:

- $t_k(b)$ is mapped to $\swarrow$ if b is matched to a red point on the left of r_k,
- $t_k(b)$ is mapped to $\downarrow$ if b is matched to r_k,
- $t_k(b)$ is mapped to $\searrow$ if b is matched to a red point on the right of r_k.

Let $w = w_1 \dots w_n$ be the word on the alphabet $\{\swarrow, \downarrow, \searrow\}$ induced by the order of the projected blue points and the map. For instance, in Figure 4.6 with $k = 3$, $w = \searrow\swarrow\swarrow\searrow\downarrow\searrow$.

Let the total order of the symbols be $\swarrow \prec \downarrow \prec \searrow$. An *inversion* in w is a pair w_i, w_j with $i < j$ and $w_j \prec w_i$. The inversions in w are in bijection with the k-observed crossing k-pairs in S. Thus, by definition, $\Phi_k(S)$ is the number of inversions in w. Lemma 4.4.4 follows from the following two observations.

(i) Any flip which is not a k-flip swaps two $\swarrow$ or two $\searrow$ in w, resulting in word w' identical to w.

(ii) Lemma 4.4.2 ensures that a crossing k-pair corresponds to an inversion in w. Thus, a k-flip exchanges the two symbols of an inversion in w, resulting in word w' with at least one inversion less than in w. $\square$

We now define $\Phi(S)$, the *potential* of S, as the sum of $\Phi_k(S)$, for k in $\{1, \dots, n\}$. The following lemma presents the key properties of Φ.

Lemma 4.4.5. *The potential Φ takes integer values from 0 to $\binom{n}{2}\frac{n+4}{3}$, and decreases by at least 2 at each flip.*

Proof. We know that Φ takes non-negative integer values by definition. By Lemma 4.4.3, an upper bound on Φ is

$$\begin{aligned}\sum_{k=1}^{n}\left(k(n+1)-k^2-1\right) &= (n+1)\sum_{k=1}^{n}k-\sum_{k=1}^{n}k^2-n\\ &= (n+1)\frac{n(n+1)}{2}-\frac{n(n+1)(2n+1)}{6}-n\\ &= \frac{n}{6}(n^2+3n-4)\\ &= \binom{n}{2}\frac{n+4}{3}.\end{aligned}$$

Finally, Lemma 4.4.4 ensures that Φ decreases by at least 2 at each flip. Indeed, a flip of a pair $\langle i,j\rangle$ is counted at least twice: once in Φ_i as an i-flip, and once in Φ_j as a j-flip. □

Theorem 4.4.1 follows from Lemma 4.4.5.

4.5 Upper Bound without Multiplicity

In this section, we prove the following theorem in the `Matching` version, yet, the proof can easily be adapted to the other versions. 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 with the same definition of the other flip. Note that there exist flip sequences (even in the `Permutation` version) using the same flip a linear number of time (Figure 4.7 [19]).

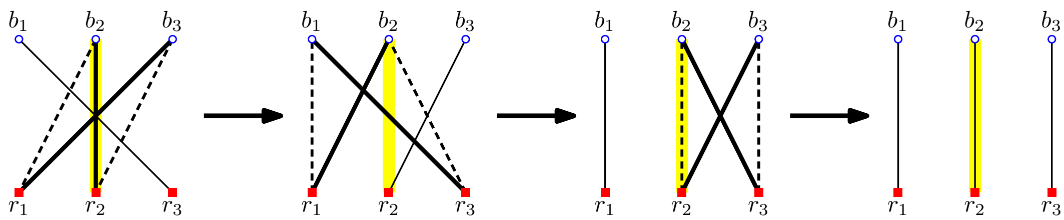

Figure 4.7: A flip sequence in the `Permutation` version where the highlighted segment r_2b_2 is removed than reinserted using 3 flips and two other segments. A similar sequence using 3 more flips and 2 more segments (not drawn) also reinsert the segment r_1b_3. It is possible to iterate this $\frac{n-1}{4}$ times, thus removing the segments r_2b_2, r_1b_3 using the same flip $\frac{n-1}{4}+1$ times. This figure appears in [19].

Theorem 4.5.1 ([30])**.** *Consider a multiset S of n segments with endpoints P.*

In the `Multigraph` *version, any untangle sequence of S has at most $O(n^{8/3})$* distinct *flips.*

The proof of Theorem 4.5.1 is based on a balancing argument from [44] and is decomposed into two lemmas that consider a flip f and two matchings S and $S' = f(S)$. Similarly to the proof of Theorem 3.1.3 [83], 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 and not on S or S'. The following lemma follows immediately from the fact that $\Lambda_L(S)$ takes integer values between 0 and $O(n^3)$.

Lemma 4.5.2. *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 4.5.2 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 4.5.3. *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.

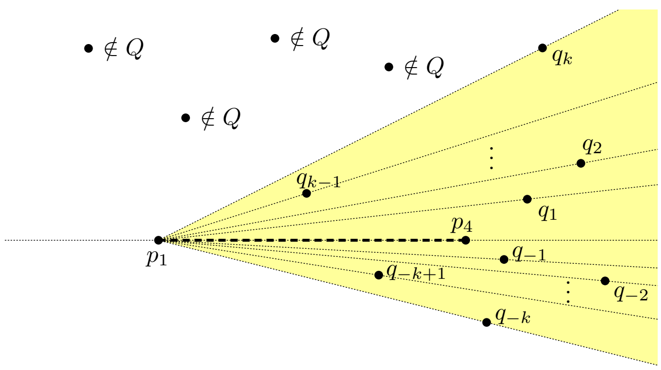

Figure 4.8: Illustration for the proof of Lemma 4.5.3.

We sweep the points in $P \setminus \{p_4\}$ by angle from the ray p_1p_4. As shown in Figure 4.8, let $q_1, \ldots, q_k$ be the first k points produced by this sweep in one direction, $q_{-1} \ldots, q_{-k}$ in the other direction and $Q = \{q_{-k} \ldots, q_{-1}, q_1, \ldots, q_k\}$. 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 $L' = \{p_1q_1, \ldots, p_1q_k\}$. 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 4.5.1 is a consequence of Lemmas 4.5.2 and 4.5.3 with $k = n^{1/3}$. S

Chapter 5

Untangling with Removal Choice

In this chapter, we first exhibit a bipartite matching on a point set in convex position such that every of its untangle sequences is long, therefore improving the lower bound on $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex}}$ by a constant factor.

Then, 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 point set versions of $\mathbf{d}^{\mathtt{R}}_{\mathtt{Multigraph}}$. Recall that such insertion 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 (the `Convex` version), followed by 1 point inside or outside the convex (the $|\mathtt{T}| = \mathtt{1}$ version), then 1 point inside and 1 outside the convex (the `Inout` version), 2 points inside the convex (the `Inin` version), and 2 points outside the convex (the `Outout` version). The last three proofs are somewhat similar but the proof of the `Inout` version is quite less tedious. As only removal choice is used, all results also apply to all the flip versions.

5.1 Lower Bound for Convex Bipartite Matchings

In this section, we prove the following linear lower bound on $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Bipartite\ Matching}}$.

Theorem 5.1.1. *Let n be an even non-negative integer. There exist a set S of n segments with endpoints P in convex position partitioned into n red points and n blue points, such that (P, S) forms a bipartite matching, and such that (in the* `Convex Bipartite Matching` *version) any untangle sequence of S is of length $\frac{3}{2}n - 2$.*

In particular, we have the following lower bound (for even n):

$$\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Bipartite\ Matching}}(n) \geq \frac{3}{2}n - 2 = \Omega(n).$$

To prove Theorem 5.1.1, we need to show that every untangle sequence starting at a given configuration (represented in Figure 5.1) is long enough. We do so by showing that every flip reduces the number of crossings by exactly 1.

We provide a convex bipartite matching which we call an *m-fence*, with $2m$ segments and $3m - 2$ crossings (Figure 5.1). Next, we give the precise definition of an m-fence, together with some useful terminology. Then, we prove Theorem 5.1.1 with three

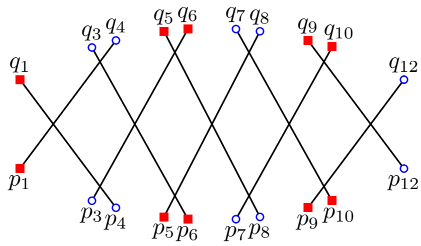

Figure 5.1: A 5-fence to lower bound $\mathbf{d}^{\mathrm{R}}(10)$.

lemmas inferring that all untangle sequences starting at an m-fence have length $3m-2$, that is, each flip reduces the number of crossings by exactly 1.

Fence. Let q_{2m+2}, q_{2m}, $q_{2m-1}, \ldots,$ q_4, q_3, q_1, p_1, p_3, $p_4, \ldots,$ p_{2m-1}, p_{2m}, p_{2m+2} be $4m$ points in convex position, ordered counter-clockwise, and with colors alternating every two points (Figure 5.1). More precisely, points p_i, q_i are red if $i \equiv 1, 2 \mod 4$ and blue otherwise. We deliberately avoid using the indices 2 and $2m+1$ to simplify the description. The segments of an m-fence are the $p_i q_{i+3}$ and the $q_i p_{i+3}$ where i is odd and varies between 1 and $2m-1$.

For $1 \leq k \leq m+1$, the *k-th column* consists of the at most 4 points with indices $2k-1$ and $2k$. We say that a convex bipartite matching with the same point set as an m-fence is a *derived m-fence* if, for all $k \in \{2, \ldots, m\}$, for all $w \in \{p, q\}$, one of the following statements holds:

1. w_{2k-1} is matched to a point of the $(k-1)$-th column, and w_{2k} is matched to a point of the $(k+1)$-th column, or

2. w_{2k-1} is matched to a point of the $(k+1)$-th column, and w_{2k} is matched to a point of the $(k-1)$-th column.

Five examples of derived m-fences are presented in Figure 5.2. Note that an m-fence is in particular a derived m-fence.

When statement 2 holds, the two segments cross. We call such a crossing an *end* crossing. Similarly, a *middle* crossing is a crossing of the form $\{p_i q_j,\ q_{i'} p_{j'}\}$, where i and i' are of the same column, and j and j' are of the same column.

Proof of Theorem 5.1.1. To prove Theorem 5.1.1, we first show with two lemmas that a flip changes a derived m-fence into another derived m-fence. Finally, we show that a flip of a derived m-fence reduces its number of crossings by exactly 1.

Lemma 5.1.2. *A crossing in a derived m-fence is either an end crossing or a middle crossing.*

Proof. The definition of a derived m-fence implies that a crossing must involve two or three consecutive columns. If exactly three columns are involved, the same definition excludes any crossing aside from the end crossings. If exactly two columns are involved, the definition again excludes any crossing aside from the middle crossings. $\square$

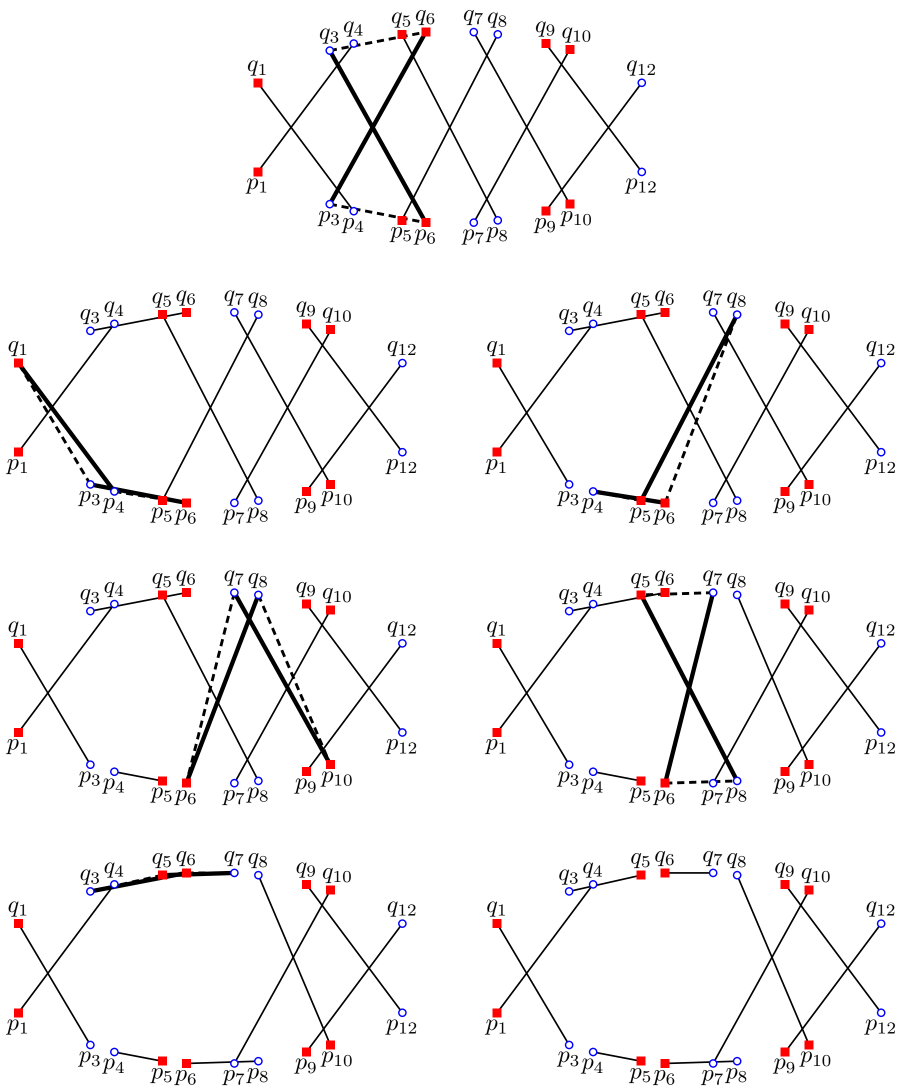

Figure 5.2: The beginning of an untangle sequence starting at a 5-fence. It is composed of derived 5-fences.

Lemma 5.1.3. *A flip changes a derived m-fence into another derived m-fence. In other words, the class of derived m-fences is closed under flip operations.*

Proof. Lemma 5.1.2 ensures that we only have the following two cases. (i) The flip of an end crossing on the w_0 side ($w_0 \in \{p, q\}$) of the k_0-th column only changes statement 2 of the definition of a derived m-fence into statement 1 for $k, r = k_0, r_0$. The statements for the other k, r are unchanged. (ii) The flip of a middle crossing simply leaves unchanged the statements for all k, r.

Figure 5.2 is actually a sequence of flips starting at an m-fence and it contains essentially all the possible cases (symmetries aside). □

Lemma 5.1.4. *A flip of a derived m-fence reduces its number of crossings by exactly* 1.

Proof. Let S be a derived m-fence. Let s_1 and s_2 be two crossing segments of S. Let s be any other segment of S. Let s_1' and s_2' be the two segments replacing s_1 and s_2 after they have been flipped, changing S into S'. We show that the number of crossings between s and s_1, s_2 is the same as between s and s'_1, s'_2, ensuring that S' has exactly 1 crossing less than S.

Let us recall that, as for any convex matching, the number of crossings cannot increase (Lemma 3.2.3 [17]). The proof of this result consists of the analysis of the five possible typical convex matchings (symmetries aside) of the three segments s_1, s_2, s (Figure 5.3). It is notable that only one of these five matchings, the one where each endpoint of s lies in between two endpoints of s_1, s_2 of the same color, corresponds to an actual decrease in the number of crossings involving s (the rightmost case in Figure 5.3).

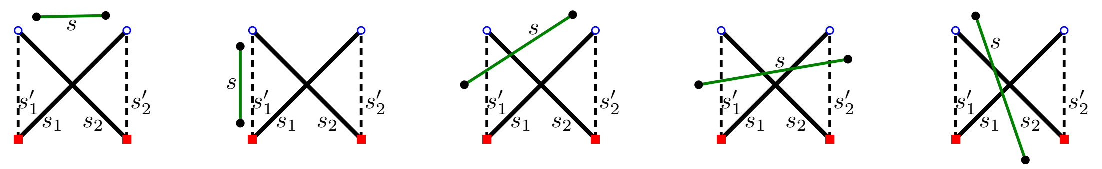

Figure 5.3: The five convex positions of s with respect to the flipping pair s_1, s_2 in the `Bipartite` version (the endpoints of the segment s may be either red or blue). This is used in the proof of Lemma 5.1.4.

This crossing-destructive case cannot occur if two endpoints of s_1, s_2 of the same color are adjacent on the convex hull. Thus, Lemma 5.1.4 holds for flips of end crossings.

If s_1, s_2 is a middle crossing, then, by definition of a derived m-fence, no segment s intersects both s_1 and s_2. Thus, the crossing-destructive case cannot occur, and Lemma 5.1.4 holds for flips of middle crossings. □

Theorem 5.1.1 follows from Lemma 5.1.3 and Lemma 5.1.4.

5.2 Upper Bound for Convex Position

Let $P = C = \{p_1, \ldots, p_{|C|}\}$ be a set of points in convex position sorted in counter-clockwise order along the convex hull boundary and consider a set of segments S with endpoints P. Given a segment $p_a p_b$ and assuming without loss of generality that $a < b$, we define the *crossing depth* $\delta_\times(p_a p_b)$ as the number of points in $p_{a+1}, \ldots, p_{b-1}$ that are an endpoint of a segment in S that crosses any other segment in S (not necessarily $p_a p_b$). This definition is also in Chapter 3. We use the crossing depth to prove an $O(n \log n)$ bound in the `Convex Multigraph` version.

Theorem 5.2.1 ([31]). *Consider a multiset S of n segments with endpoints $P = C$ in convex position.*

In the `Convex Multigraph` *version, there exists a removal strategy such that for any insertion strategy, the resulting untangle sequence of S is of length $O(n \log |C|) = O(n \log n)$.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Multigraph}}(n) = O(n \log |C|) = O(n \log n).$$

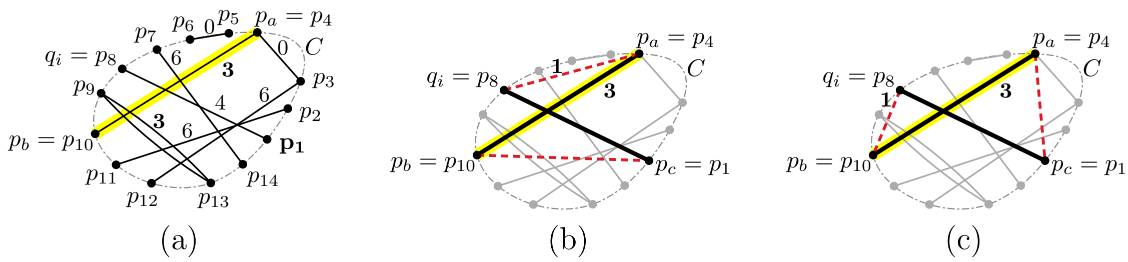

Figure 5.4: Proof of Theorem 5.2.1. (a) The segments of a convex multigraph are labeled with the crossing depth. (b,c) Two possible pairs of inserted segments, with one segment of the pair having crossing depth $\lfloor \frac{3}{2} \rfloor = 1$.

Proof. We repeat the following procedure until there are no more crossings. Let $p_a p_b \in S$ be a segment *with crossings* (hence, crossing depth at least one) and $a < b$ minimizing $\delta_\times(p_a p_b)$ (Figure 5.4(a)). Let $q_1, \ldots, q_{\delta_\times(p_a p_b)}$ be the points defining $\delta_\times(p_a p_b)$ in order and let $i = \lceil \delta_\times(p_a p_b)/2 \rceil$. Since $p_a p_b$ has minimum crossing depth, the point q_i is the endpoint of segment $q_i p_c$ that crosses $p_a p_b$. When flipping $q_i p_c$ and $p_a p_b$, 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 $\delta_\times(p_a p_b)$ (Figure 5.4(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. $\square$

5.3 Upper Bound for Convex Trees

In this section, we prove the following linear upper bound in the `Convex Tree` version.

Theorem 5.3.1. *Consider a set S of n segments with endpoints $P = C$ in convex position such that (P, S) forms a tree.*

In the `Convex Tree` *version, for $n \geq 3$ there exists an untangle sequence of S of length at most $3n - 8 = O(n)$.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\mathtt{R}}_{\texttt{Convex Tree}}(n) \leq 3n - 8 = O(n) \qquad \textit{for } n \geq 3.$$

The proof of Theorem 5.3.1 has the same structure as the proof of Theorem 3.2.7. It relies on the following two lemmas.

Lemma 5.3.2. *Consider a set S of segments with endpoints P such that (P, S) forms a tree.*

If there exists a segment $p_a p_b \in S$ and a point $q \in P$ of degree at least 2 such that all the segments of S incident to q cross the segment $p_a p_b$, then there exists a point $p_c \in P$ such that $q p_c \in S$ and such that the flip (in the `Tree` *version) removing the segments $p_a p_b, q p_c$ inserts the segment $p_a q$.*

Proof. Let p_d be a point in P such that $q p_d \in S$ and such that the flip (in the `Tree` version) removing the segments $p_a p_b, q p_d$ does not insert the segment $p_a q$. By definition of a flip in the `Tree` version, the unique path $u_{q\to p_a}$ connecting q to p_a in the tree (P, S) (before the flip) includes either none or both of the segments p_ap_b, qp_d (to avoid forming a cycle with the inserted segments).

Case 1: $u_{q\to p_a}$ includes none. In the case where $u_{q\to p_a}$ includes none of the segments p_ap_b, qp_d, then let qp_c be the first segment of $u_{q\to p_a}$ (starting at the point q). This segment qp_c crosses the segment p_ap_b by hypothesis. The path $u_{q\to p_a}$ includes the segment qp_c but not the segment p_ap_b. Thus, by definition of a flip in the `Tree` version, the flip removing the segments p_ap_b, qp_c inserts the segment p_aq.

Case 2: $u_{q\to p_a}$ includes both. In the case where $u_{q\to p_a}$ includes both of the segments p_ap_b, qp_d, then let qp_c be any other segment adjacent to q. Such a segment exists since the point q has degree at least 2. The path $u_{q\to p_a}$ includes the segment p_ap_b but not the segment qp_c. Thus, by definition of a flip in the `Tree` version, the flip removing the segments p_ap_b, qp_c inserts the segment p_aq. □

Lemma 5.3.3. *Any set S of n segments with endpoints P in convex position such that (P, S) forms a tree admits an uncrossable segment $s \in S$ after at most 2 flips (in the* `Convex Tree` *version).*

Proof. We assume that S contains no segment of null crossing depth. Let $p_1, \dots, p_{|P|}$ be the points of P sorted in counterclockwise order along the convex hull boundary. Let $p_ap_b \in S$ $(a < b)$ be a segment *with crossings* (hence, of crossing depth at least one) minimizing $\delta_\times(p_ap_b)$. Let $q_1, \dots, q_{\delta_\times(p_ap_b)}$ be the points defining $\delta_\times(p_ap_b)$ sorted in counterclockwise order along the convex hull boundary. Since p_ap_b has minimum positive crossing depth, and since we assume that S contains no segment of null crossing depth, all the segments with at least one endpoint in $q_1, \dots, q_{\delta_\times(p_ap_b)}$ cross the segment p_ap_b. We use 1 or 2 flips to insert one of the segments $q_1p_a, q_{\delta_\times(p_ap_b)}p_b$.

Case 1: using 1 flip. This case is defined as all the situations where one flip is enough to insert one of the segments $q_1p_a, q_{\delta_\times(p_ap_b)}p_b$. Formally, there is nothing to prove. Note that, by Lemma 5.3.2, Case 1 covers all the situations where at least one of $q_1, q_{\delta_\times(p_ap_b)}$ has degree more than 1.

Case 2: using 2 flips. This case is defined as the complement of Case 1. Therefore, both q_1 and $q_{\delta_\times(p_ap_b)}$ have degree exactly 1. Let $p_c \in P$ (respectively $p_d \in P$) be the other endpoint of the segment adjacent to q_1 (respectively to $q_{\delta_\times(p_ap_b)}$). We choose to remove the segments q_1p_c, p_ap_b. By definition of Case 2, this flip inserts the segments q_1p_b, p_ap_c. We now choose to remove the segments $q_1p_b, q_{\delta_\times(p_ap_b)}p_d$. As the segment $q_1q_{\delta_\times(p_ap_b)}$ would not preserve the `Tree` property if inserted, this flip inserts the segment $q_{\delta_\times(p_ap_b)}p_b$ which has null crossing depth. □

Proof of Theorem 3.2.7. For induction purpose, let $f(n)$ be the length of the untangle sequence of S in the statement of Theorem 3.2.7 for n segments.

If $n = 3$, then $f(n) = 1 = 3n - 8$.

If $n \geq 4$, and if $f(n-1) \leq 3(n-1) - 8$, then, by Lemma 5.3.3, using at most 2 flips ensures that there exists at least one segment s in S contained in the boundary of the convex hull of P. We invoke the induction hypothesis on the tree with $n-1$ segments obtained from S by contracting the segment s to a point. The resulting flip sequence also transforms S into a tree where only the two segments adjacent to s may cross. A final flip is enough to fully untangle S. The total number of flips use to untangle S is

$$f(n) \leq 3 + f(n-1) \leq 3 + 3(n-1) - 8 = 3n - 8,$$

concluding the induction. $\square$

5.4 Upper Bound for One Point Inside or Outside a Convex

In this section, we prove Theorem 5.4.2, in which we assume that the multiset of segments to be untangled has no crossing pair of CC-segments (the segments with both endpoints in C where C is a subset of the endpoints which is in convex position), possibly after having untangled the CC-segments using one of the theorems of the `Convex` version. In the `Cycle` version (Theorem 3.2.7 [72] and Theorem 3.2.9 [85]) and in the `Bipartite Matching` version (Theorem 3.2.10 [18]), the preprocessing to untangle CC-segments takes $O(n)$ flips. However, in the `Tree` version (Theorem 3.2.11 [18]) and in general (Theorem 5.2.1 [31]), the best bound known is $O(n \log n)$. We first state a lemma used to prove Theorem 5.4.2.

Lemma 5.4.1 ([31]). *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.*

In the $|\mathtt{T}| = 1$ `Multigraph` *version, there exists a removal strategy such that, for any insertion strategy, the resulting untangle sequence of $S \cup \{s\}$ is of length $O(n)$.*

Proof. Iteratively flip the segment qp_1 with the segment $p_2p_3 \in S$ crossing qp_1 the farthest from q. This flip inserts a CC-segment p_1p_2, which is impossible to flip again, because the line p_1p_2 is crossing free. The flip does not create any crossing between CC-segments. $\square$

We are now ready to state and prove the theorem.

Theorem 5.4.2 ([31]). *Consider a multiset S of n segments with endpoints P partitioned into $P = C \cup T$ where C is in convex position, and $T = \{q\}$, and such that S has no crossing pair of CC-segments (the segments with both endpoints in C), possibly after having untangled the CC-segments using one of the theorems of the* `Convex` *version. We define the parameter t as the sum of the degrees of the points in T.*

In the $|\mathtt{T}| = 1$ `Multigraph` *version, there exists a removal strategy such that for any insertion strategy, the resulting untangle sequence of S is of length $O(tn)$.*

In particular (using Theorem 5.2.1 to preprocess CC-segments), we have the following upper bound:

$$\mathbf{d}^{\mathtt{R}}_{|\mathtt{T}|=1\ \mathtt{Multigraph}}(n,t) = O(n \log |C| + tn) = O(n \log n + tn).$$

Proof. For each segment s with endpoint q with crossing, we apply Lemma 5.4.1 to s and the CC-segments 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. Let ℓ be the line containing s.

Case 1: If ℓ is crossing free, then ℓ splits (see Lemma 2.6.1 [19, 31]) 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 2.6.1 [19, 31]).

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 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, and, therefore, that one of the removed segment already has q as an endpoint.

Therefore, we need at most n flips for each of the t segments incident to q. □

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

Given an endpoint p, let $d^\circ(p)$ denote the degree of p, that is, the number of segments incident to p. The following lemma is used to prove Theorem 5.5.2.

Lemma 5.5.1. *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., $t = d^\circ(q') + 1$.*

In the `Inout Multigraph` *version, there exists a removal strategy such that, for any insertion strategy, the resulting flip sequence starting at S is of length $O(tn)$ and ends with a multiset of segments where with all crossings (if any) are on the segment qq' (if $qq' \notin S$ then there are no crossings).*

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 5.5).

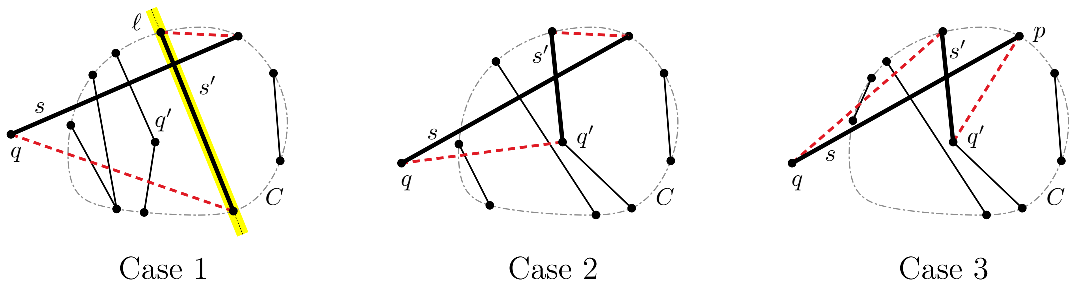

Figure 5.5: The three cases in the proof of Lemma 5.5.1.

Case 1 ($\mathbf{CT} \times \mathbf{CC}$). 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 the splitting lemma, i.e., Lemma 2.6.1 [19, 31]) using $O(n)$ and $O(tn)$ flips (Theorem 5.4.2). 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 ($\mathbf{CT} \times \mathbf{CT} \to \mathbf{CC}, \mathbf{TT}$). 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 ($\mathbf{CT} \times \mathbf{CT} \to \mathbf{CT}, \mathbf{CT}$). 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. □

We are now ready to prove the theorem.

Theorem 5.5.2 ([31]). *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. We define the parameter t as the sum of the degrees of the points in T, i.e., $t = d^\circ(q) + d^\circ(q')$.*

In the `Inout Multigraph` *version, there exists a removal strategy such that for any insertion strategy, the resulting untangle sequence of S is of length $O(d^\circ(q)d^\circ(q')n + d_{conv}(n)) = O(t^2 n + d_{conv}(n))$, where $d_{conv}(n)$ is the number of flips to untangle any multiset of at most n segments with endpoints in convex position.*

In particular (using Theorem 5.2.1 for $d_{conv}(n)$), we have the following upper bound:

$$\mathbf{d}^{\mathrm{R}}_{\texttt{Inout Multigraph}}(n,t) = O(d^\circ(q)d^\circ(q')n + n\log n) = O(t^2 n + n \log n).$$

Proof. The untangle sequence contains four phases.

Phase 1 ($\mathbf{CC} \times \mathbf{CC}$). In this phase, we remove all crossings between pairs of CC-segments using $d_{\mathrm{conv}}(n)$ flips. Throughout all the phases, the invariant that no pair of CC-segments crosses is preserved.

Phase 2 ($\mathbf{Cq'} \times \mathbf{CC}$). 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 5.4.2.

Phase 3 ($\mathbf{Cq}$). 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'. The splitting lemma (Lemma 2.6.1 [19, 31]) allows us to remove the crossings in each CT-segment s incident to q independently, which we do using $O(d^\circ(q')n)$ flips using Lemma 5.5.1. As there are $d^\circ(q)$ CT-segments adjacent to q, the total number of flips is $O(d^\circ(q)d^\circ(q')n) = O(t^2 n)$.

Phase 4 (CC × 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 5.4.2 using $O(tn)$ flips.

After the $d_{\text{conv}}(n)$ flips in Phase 1, the number of flips is dominated by Phase 3 with $O(d^\circ(q)d^\circ(q')n) = O(t^2n)$ flips. □

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 5.5.1 never happens, and the bound in Theorem 5.5.2 decreases to $O(d_{\text{conv}}(n) + tn)$.

5.6 Upper Bound for Two Points Inside a Convex

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

Theorem 5.6.1 ([31]). *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 and q' are inside the convex hull of C. We define the parameter t as the sum of the degrees of the points in T.*

In the `Inin Multigraph` *version, there exists a removal strategy such that for any insertion strategy, the resulting untangle sequence of S is of length $O(tn + d_{conv}(n))$, where $d_{conv}(n)$ is the number of flips to untangle any multiset of at most n segments with endpoints in convex position.*

In particular (using Theorem 5.2.1 for $d_{conv}(n)$), we have the following upper bound:

$$\mathbf{d}^{\mathtt{R}}_{\mathtt{Inin\ Multigraph}}(n, t) = O(tn + n \log n).$$

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 × CT). In this phase, we remove all crossings between pairs of CT-segments using $O(d_{\text{conv}}(t)) = O(d_{\text{conv}}(n))$ flips. We separately solve two convex sub-problems defined by the CT-segments, one on each side of the line qq'.

Phase 2 (CC × CC). In this phase, we remove all crossings between pairs of CC-segments using $O(d_{\text{conv}}(n))$ 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 crosses.

Phase 3 (CT × 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 and non-central CC-segments.

Given a non-central CC-segment pp', let the *out-depth* of a 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. At the end of each step the two following invariants are preserved. (i) No pair of CC-segments crosses. (ii) No pair of CT-segments crosses.

At each step, we choose to flip the 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 5.6(a) and Figure 5.7(a)). One of the possibly inserted pairs may contain a CT-segment s that crosses another CT-segment s', violating the invariant (ii) (Figure 5.6(b) and Figure 5.7(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 5.6(c) and Figure 5.7(c)) or a CC-segment and the segment qq' (Figure 5.6(d) and Figure 5.7(d)). The analysis is divided in two main cases.

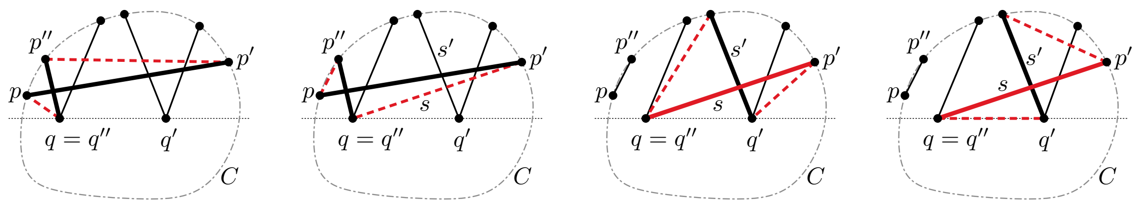

(a) (b) (c) (d)

Figure 5.6: Theorem 5.6.1, Phase 3 when pp' is an outermost segment.

If pp' is an outermost CC-segment (see Figure 5.6), then case analysis shows that the two invariants are preserved and χ decreases.

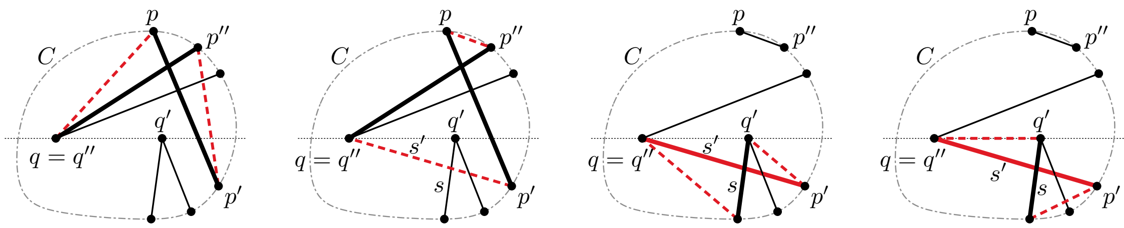

(a) (b) (c) (d)

Figure 5.7: Theorem 5.6.1, Phase 3 when pp' is a peripheral segment.

If pp' is a peripheral CC-segment (see Figure 5.7), 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 5.7(d)). Notice that the number of times the TT-segment qq' is inserted is $O(t)$, which bounds the total increase by $O(t^2)$.

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 ($\mathbf{CT} \times \mathbf{CC}_{\mathbf{central}}$). 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, ignoring the TT-segments. This phase ends with crossings only between qq' and central CC-segments.

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

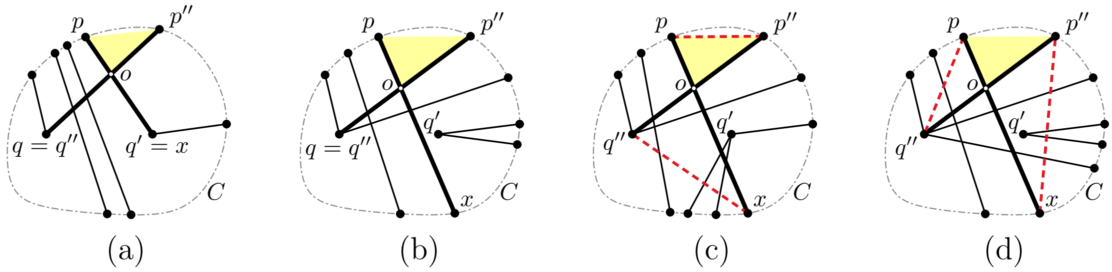

Figure 5.8: Theorem 5.6.1, Phase 4. (a) A pair of CT-segments with an ear. (b) A CC-segment and a CT-segment with an ear. (c) Flipping an ear that produces crossing pairs of CT-segments. (d) Flipping an ear that inserts a non-central CC-segment with crossings.

triangle $pp''o$ intersects no segment of S (see Figure 5.8(a) and 5.8(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, they may be produced in the flip (Figure 5.8(c)). If the flip inserts a non-central CC-segment which crosses some CT-segments (Figure 5.8(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 plus the number of crossings between CT-segments. 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 the splitting lemma, i.e., Lemma 2.6.1 [19, 31]). As the number of CC-segments 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. The endpoints of the segments with crossings are in convex position and all other endpoints are outside their convex hull. Hence, by the splitting lemma (Lemma 2.6.1 [19, 31]), it is possible to obtain a crossing-free multigraph using $O(d_{\text{conv}}(n))$ flips. $\square$

5.7 Upper Bound for Two Points Outside a Convex

In this section, we prove a theorem with a bound that is exponential in t, which makes it of little interest for large t. Notice, however, that in matchings $t \leq 2$, in a TSP tour $t \leq 4$, and in a binary tree $t \leq 6$. Also notice that the definition of t is different from other theorems (here TT-segments are counted twice). Both definitions are equivalent up to a factor of 2, but since t appears in the exponent, they are not exchangeable.

Theorem 5.7.1 ([31]). *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 and q' are outside the convex hull of C. We define the parameter t as the sum of the degrees of the points in T.*

In the `Outout Multigraph` *version, there exists a removal strategy such that for any insertion strategy, the resulting untangle sequence of S is of length $O(2^t d_{conv}(n))$, where $d_{conv}(n)$ is the number of flips to untangle any multiset of at most n segments with endpoints in convex position.*

In particular (using Theorem 5.2.1 for $d_{conv}(n)$), we have the following upper bound:

$$\mathbf{d}^{\mathtt{R}}_{\mathtt{Outout\ Multigraph}}(n, t) = O(2^t n \log n).$$

Proof. Throughout this proof, we partition the TT-segments respectively the CT-segments into two types: *TT_{in}-segment* and *CT_{in}-segment* if it intersects the interior of the convex hull of C and *TT_{out}-segment* and *CT_{out}-segment* otherwise. Let $f(t)$ be the number of flips to untangle a multiset S as in the statement of the theorem. The proof proceeds by induction. The base case is $t = 0$, when $f(0) \leq d_{\text{conv}}(n)$ by definition of $d_{\text{conv}}(n)$.

Next, we show how to bound $f(t)$ for $t > 0$, but first we need some definitions. A *T-segment* is a segment of S with at least one of its endpoints in T. A line ℓ is a *T-splitter* if ℓ splits S and either ℓ contains a T-segment or there are T-segments on both sides of ℓ. We abusively say that a segment s is a T-splitter if the line containing s is a T-splitter. A T-splitter is useful because we can apply the splitting lemma (Lemma 2.6.1 [19, 31]) and solve sub-problems with a lower value of t by induction.

Phase 1: untangle all but one segment by induction. We remove an arbitrary CT-segment or TT-segment s from S. We then use induction to untangle S using $f(t-1)$ flips and insert the segment s back in S afterwards. Notice that all crossings are now on s.

Phase 2.1: apply induction if possible. If S admits a T-splitter ℓ, then we apply the splitting lemma (Lemma 2.6.1 [19, 31]) to solve each side of ℓ independently using induction.

If S has a crossing-free TT_{out}-segment qq' such that the line qq' is not crossing free, then qq' is uncrossable, and we remove qq' from S and untangle S by induction. Similarly, in the case where $T = \{q, q'\}$ and where qq' is a TT_{in}-segment, if S has a CT_{out}-segment, say pq, then pq is uncrossable, and we remove pq from S and untangle S by induction.

In all the three cases of Phase 2.1 we get $f(t) \leq f(t-1) + f(t_1) + f(t_2)$, where $t_1 + t_2 \leq t$ and $t_1, t_2 \geq 1$.

Phase 2.2: split after one flip. If S contains no T-splitter and if s is a TT-segment, then there remains no CT-segment in S (as every CT-segment shares an endpoint with the TT-segment s that contains all crossings), and s crosses a CC-segment s'. A crossing-free CT-segment would either be a CT_{in}-segment, hence a T-splitter, or a CT_{out}-segment and, hence uncrossable and removed by one of the induction cases of Phase 2.1.

The segment s' becomes a T-splitter after flipping s with s', and we invoke induction. By the splitting lemma (Lemma 2.6.1 [19, 31]), we get in this case $f(t) \leq f(t-1)+1+f(t_1)+f(t_2)$, where $t_1+t_2 \leq t$ and $t_1, t_2 \geq 1$.

Phase 2.3: split after $O(n)$ flips. In this case, S contains no T-splitter and s is a CT-segment, say with q as its endpoint in T. While s', the segment of S that crosses s the farthest away from q, is a CC-segment, we flip s and s' and we set s to be the newly inserted CT-segment incident to q. By Lemma 5.4.1, at most n flips are performed in this loop.

At the end of the loop, either s is crossing free, or s' is a CT-segment, say with q' as its endpoint in T. Then, we also flip s and s'.

Insertion case 1: If two CT-segments are inserted, then, either one of them is uncrossable (this is the case if s' is a CT_{out}-segment), or s' is now a T-splitter (recall that if qq' is a TT_{in}-segment, then all the CT_{out}-segments have been removed at Phase 2.1.).

Insertion case 2: If the TT-segment qq' is inserted, then the inserted CC-segment is crossing free (as in the proof of Lemma 5.4.1), and, if qq' is not already crossing free, we flip qq' with any segment, say pp'.

Next, we split S as follows. Among the CT_{in}-segments of S which are on the upper (respectively lower) side of the line qq', consider the one whose endpoint p_{upper} (respectively p_{lower}) in C is the closest to the line qq'. The segments of S are either inside or outside the convex quadrilateral $qp_{\text{lower}}q'p_{\text{upper}}$, and we know that only the segments inside may have crossings. By the splitting lemma (Lemma 2.6.1 [19, 31]), we remove from S all the segments outside $qp_{\text{lower}}q'p_{\text{upper}}$. Recall that, in our case, qq' is a TT_{in}-segment, and all the CT_{out}-segments have been removed at Phase 2.1. The line pp' is finally a T-splitter. Again, by the splitting lemma (Lemma 2.6.1 [19, 31]), we get in this case $f(t) \leq f(t-1)+n+2+f(t_1)+f(t_2)$, where $t_1+t_2 \leq t$ and $t_1, t_2 \geq 1$.

The last bound on $f(t)$ dominates the recurrence. Using that $f(t_1)+f(t_2) \leq f(t-1)+f(1)$ and $t<n$ we get

$$f(t) \leq f(t-1)+n+2+f(t_1)+f(t_2) \leq O(n)+2f(t-1),$$

which solves to $f(t)=O(2^t d_{\text{conv}}(n))$ as claimed. □

5.8 Upper Bound for Red-on-a-Line Matchings

In this section, we prove the following upper bound.

Theorem 5.8.1 ([33]). *Consider a set S of n segments with endpoints P partitioned into n red points and n blue points such that the red points lie on the x-axis, such that the blue points lie above the x-axis, and such that (P, S) forms a bipartite matching.*

In the `Redonaline Matching` *version, there exists an untangle sequence of S of length at most $\binom{n}{2} = O(n^2)$.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\mathrm{R}}_{\texttt{Redonaline Matching}}(n) \leq \binom{n}{2} = O(n^2).$$

The proof consists of the analysis of the number of flips performed by the recursive algorithm described next. This analysis is based on a novel approach called *state tracking*. State tracking is in fact not specific to the red-on-a-line case, which is why Lemma 5.8.2 is stated and proven in the non-bipartite setting. Lemma 5.8.2 is then used in the red-on-a-line case to prove Lemma 5.8.3, which in turn is used to prove Theorem 5.8.1. Lemma 5.8.2 also provides an alternative proof of the well-known Theorem 3.2.2 [17], which we present at the end of this section.

Throughout, we assume general position (no two blue points with the same y-coordinate). Let the *top* segment of a red-on-a-line matching be the segment with the topmost blue endpoint (Figure 5.9(a)).

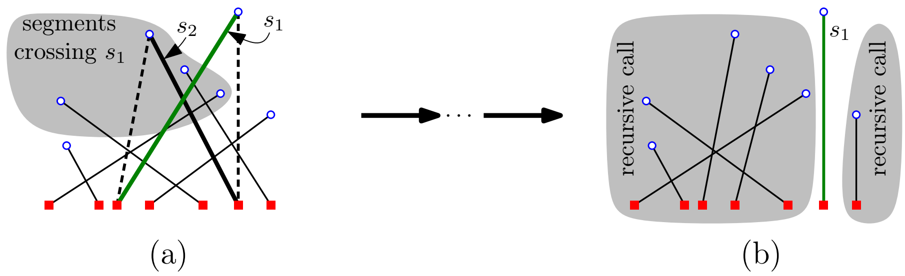

Figure 5.9: (a) A red-on-a-line matching with s_1 as the top segment. (b) The matching just before the first recursive calls of the algorithm, where s_1 is uncrossable.

Algorithm. While the top segment s_1 of the matching crosses another segment s_2, we flip s_1 and s_2. If multiple segments cross s_1, then we choose s_2 as the top segment among the segments crossing s_1.

The previous loop stops when the top segment s_1 has no crossings. At this point, we have that s_1 splits (Lemma 2.6.1 [19, 31]) the matching into at most two non-empty submatchings, one to each side of s_1. We recursively call the algorithm on these submatchings (Figure 5.9(b)).

Flip Complexity. The analysis of the number of flips performed by the algorithm stems from the following observations. We define three possible *states* for a pair of segments (Figure 5.10).

- State **X**: the segments are crossing.
- State **H**: the segments are not crossing and their endpoints are in convex position.
- State **T**: the endpoints are not in convex position.

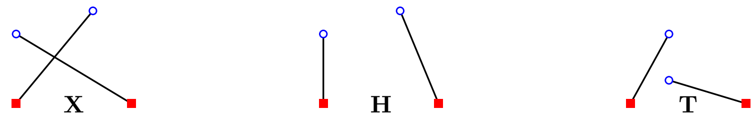

Figure 5.10: The three different states of pairs of segments.

In the convex case, there are no **T**-states and a flip increases the number of **H**-pairs by at least 1, and decreases the number of **X**-pairs as well. Hence, counting either **X** or **H**-pairs yields the $\binom{n}{2}$ upper bound on $\mathbf{d}^{\emptyset}_{\texttt{Convex Multigraph}}(n)$ (see below the alternative proof of Theorem 3.2.2 [17]). However, when the points are not in convex position, counting **H** and **X**-pairs is fundamentally different. We will see that counting **H**-pairs is more useful to prove the desired bounds.

When the points are not in convex position, a flip may decrease the number of **H**-pairs. Figure 5.11 shows two such situations where flipping s_1, s_2 does not increase the number of **H**-pairs. There is one **H**-pair involving segment s before the flip, and none after the flip. Notice that, if we added multiple segments close to s, the number of **H**-pairs would actually decrease. However, the algorithm avoids these situations by choosing to flip top segments. The full proof involves state tracking, a novel approach to analyze flip sequences, which is described next.

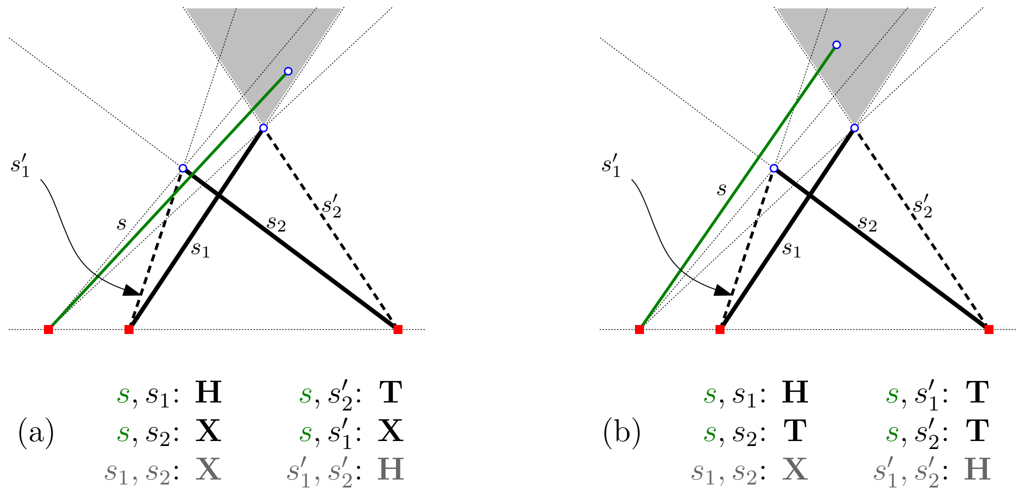

Figure 5.11: Two cases where flipping s_1, s_2 does not increase the number of **H**-pairs. The upper cone of s_1, s'_2 is shaded.

State Tracking. We have $\binom{n}{2}$ pairs of segments before and after a flip. Each pair has an associated state. However, since two segments change in the matchings, there is no clear correspondence between the state of each pair before and after the flip. State tracking establishes this correspondence by making choices of which pair of segments in the initial matching corresponds to which pair of segments in the resulting matching. These choices are performed deliberately to obtain certain state transitions instead of others and prove the desired bounds.

The following notations will be used throughout the rest of this section and are summarized in Figure 5.12. Let r_1, r_2 be two red points and b_1, b_2 be two blue points.

Let s_1, s_2, s'_1, s'_2 be the following four segments respectively: r_1b_1, r_2b_2, r_1b_2, r_2b_1. We consider a flip that replaces the pair of segments s_1, s_2 by s'_1, s'_2. Let S denote the matching before the flip and S' denote the resulting matching after the flip.

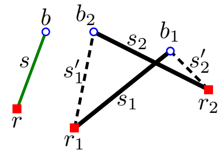

Figure 5.12: Notations for a generic flip and for a variable segment s.

We order the $\binom{n}{2}$ pairs of segments of S in a column vector. There are three types of pairs of segments in S with respect to the flip: the *unaffected* pairs (involving neither s_1 nor s_2), the *flipping* pair s_1, s_2, and the *affected* pairs (involving exactly one of s_1 or s_2). We choose the new order of the $\binom{n}{2}$ pairs of segments of S' in a way that satisfies the following properties with respect to the previous vector. The unaffected pairs keep the same indices. The pair s'_1, s'_2 gets the index of s_1, s_2. Next, we describe the remaining indices.

Let s be a segment of S distinct from s_1 and s_2. Let r and b be the red and blue endpoints of s. Let i_1 and i_2 be the indices of s, s_1 and s, s_2, and let $\mathbf{S}_1$ and $\mathbf{S}_2$ be their respective states. Let $\mathbf{S}'_1$ and $\mathbf{S}'_2$ be the respective states of s, s'_1 and s, s'_2. We restrict our choice to the following two options:

- index s, s'_1 with i_1, and s, s'_2 with i_2, or
- index s, s'_1 with i_2, and s, s'_2 with i_1.

We call such a choice a *tracking choice*. We say that a pair of segments in S *turns into* a pair in S' when they have the same index. We denote $\mathbf{S} \to \mathbf{S}'$ to specify that the pairs of segments with a given index go from the state $\mathbf{S}$ to the state $\mathbf{S}'$. In the following, we use $\mathbf{S}_1\mathbf{S}_2 \to \mathbf{S}'_1\mathbf{S}'_2$ as a shorthand notation to say that we have the two following tracking choices: either $\mathbf{S}_1 \to \mathbf{S}'_1$ and $\mathbf{S}_2 \to \mathbf{S}'_2$ or $\mathbf{S}_1 \to \mathbf{S}'_2$ and $\mathbf{S}_2 \to \mathbf{S}'_1$.

There are 3^2 possible such *transitions* $\mathbf{S} \to \mathbf{S}'$. Yet, the next two lemmas ensure that some transitions can be ruled out by tracking choices. Lemma 5.8.2 actually holds for any (possibly non-bipartite) matching, while Lemma 5.8.3 is specific to the red-on-a-line case. Both lemmas are proved analyzing the tracking choices of each possible position of a segment s relatively to the flipping pair.

Lemma 5.8.2. *There always exists a tracking choice avoiding the* $\mathbf{H} \to \mathbf{X}$ *transition.*

Proof. There clearly exists a tracking choice avoiding the $\mathbf{H} \to \mathbf{X}$ transition unless we have either a transition (i) $\mathbf{HH} \to \mathbf{XS}$ or (ii) $\mathbf{HS} \to \mathbf{XX}$, where $\mathbf{S} \in \{\mathbf{X}, \mathbf{H}, \mathbf{T}\}$. We show that these two cases are not possible.

(i) $\mathbf{HH} \to \mathbf{XS}$: If both the pairs s, s_1 and s, s_2 are $\mathbf{H}$ while at least one of the two pairs s, s'_1 and s, s'_2 is $\mathbf{X}$, then the final $\mathbf{X}$ state implies that s crosses s_1 or s_2, which contradicts the two initial $\mathbf{H}$ states.

(ii) $\mathbf{HS} \to \mathbf{XX}$: If one of the two pairs s, s_1 and s, s_2 is $\mathbf{H}$ while both pairs s, s'_1 and s, s'_2 are $\mathbf{X}$, then the two final $\mathbf{X}$ states imply that s crosses s'_1 and s'_2. It follows that s also crosses s_1 and s_2, which is again a contradiction. □

State Tracking in the Red-on-a-Line Case. Figure 5.13 summarizes the notations for a generic red-on-a-line flip and an variable segment s. Figures 5.14, 5.15, 5.16, and 5.17 then provide "maps" of essentially all the possible situations of tracking choices in the red-on-a-line case. These figures are used to prove the next lemma.

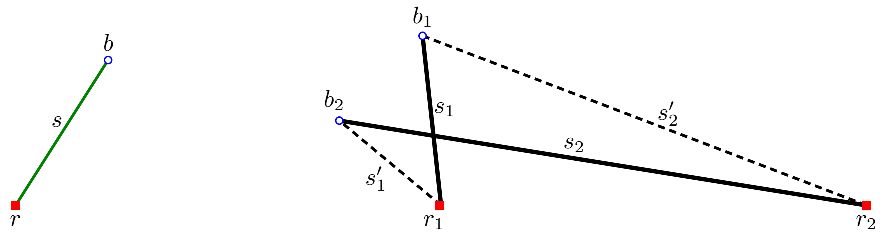

Figure 5.13: Notations used in Figures 5.14, 5.15, 5.16, and 5.17 for a generic red-on-a-line flip and an variable segment s.

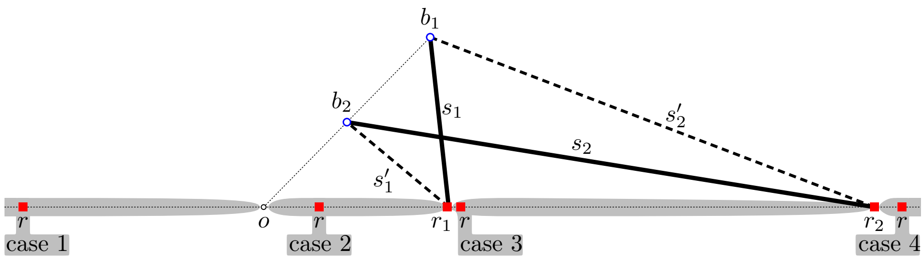

Figure 5.14: The four possible cases for the position of r.

Figures 5.15, 5.16, and 5.17 are generated by a brute force computation of the states $\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_1', \mathbf{S}_2'$ of the four pairs s, s_1, s, s_2, s, s_1', s, s_2' for each position case for r (Figure 5.14) and for each position case for b (in Figures 5.15, 5.16, and 5.17, each cell of the arrangement of lines corresponds to a position case for b). In the following, we make sure that no case is forgotten.

We assume, without loss of generality, that r_1 is on the left of r_2, and that b_1 is higher than b_2. Let o be the intersection between the line b_1b_2 and the red-point line. There are, indeed, four possible open intervals for the position of r on the red-point line: $]-\infty, o[$, $]o, r_1[$, $]r_1, r_2[$, and $]r_2, \infty[$ (Figure 5.14). This yields four cases, respectively. We do not explicitly describe case 4 as it is similar to case 2. Indeed, case 2 and case 4 map to each other by exchanging the labels of r_1 and r_2, as well as b_1 and b_2. The fact that the point o is still on the left of r_1 and r_2 is not a problem since we are studying incidence proprieties. Another way to see it, is to consider the projective plane.

As we have assumed the blue points to lie in the upper half-plane, these four cases branch into further sub-cases. However, no-loss-of-generality assumptions and symmetries simplify the analysis. Without loss of generality, we first assume that the lines r_1b_2 and r_2b_1 intersect in the upper half-plane, as it will only generate more cells to the upper part of the arrangement of lines.

Second, we examine case 3. Let o_1 be the intersection of the lines rb_2 and r_2b_1 (see Figure 5.17), and o_2 be the intersection of the lines rb_1 and r_1b_2. Case 3 decomposes into:

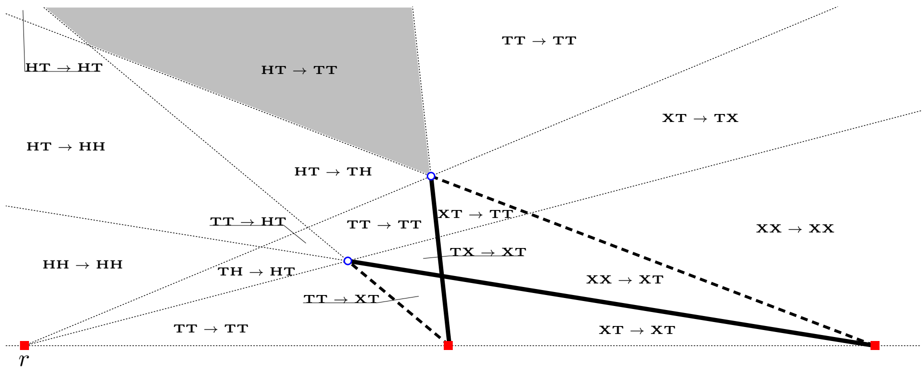

Figure 5.15: The case 1 "map" of all the possible red-on-a-line tracking choices. Tracking choices cannot avoid the transition **H** → **T** in the shaded region.

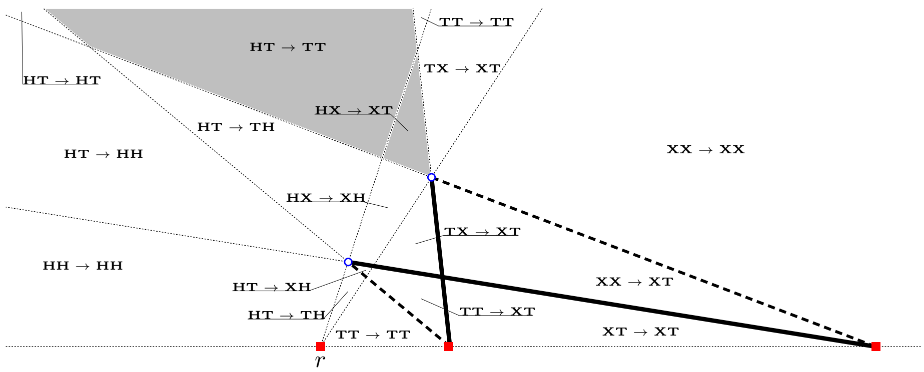

Figure 5.16: The case 2 "map" of all the possible red-on-a-line tracking choices. Tracking choices cannot avoid the transition **H** → **T** in the two shaded regions.

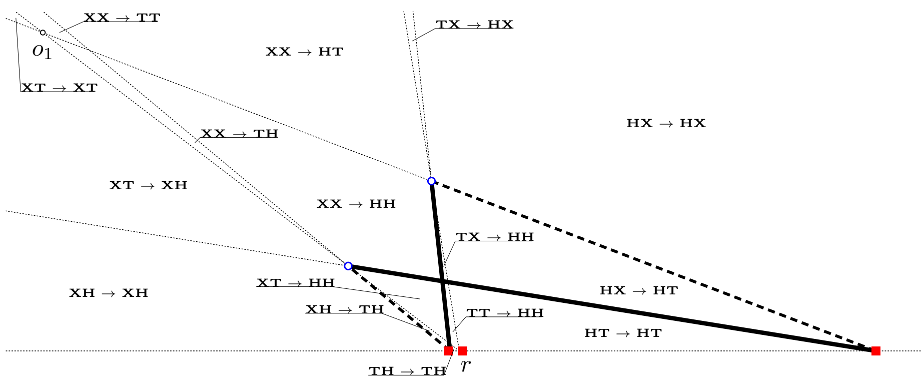

Figure 5.17: The case 3.1 "map" of all the possible red-on-a-line tracking choices.

- case 3.1 where o_1 lies in the upper half-plane and o_2 in the lower,
- case 3.2 where both o_1 and o_2 lie in the upper half-plane,
- case 3.3 where o_1 lies in the lower half-plane and o_2 in the upper, and
- case 3.4 where both o_1 and o_2 lie in the lower half-plane.

Cases 3.1 and 3.3 are similar, while case 3.2 is just a superposition of both of them. More precisely, when compared to case 3.4, the extra cell of the arrangement generated by case 3.1 (the cell in the top left corner of Figure 5.17) corresponds to the possible tracking choices summarized by the notation $\mathbf{XT} \to \mathbf{XT}$. Similarly, the extra cell generated by case 3.3 corresponds to $\mathbf{TX} \to \mathbf{TX}$. The two extra cells generated by case 3.2 are the same as the two previous ones. We thus assume case 3.1 (as it is easier to draw in our setting) without loss of generality. All these assumptions made, the remaining cases now correspond to Figures 5.15, 5.16, and 5.17.

The next lemma is similar to Lemma 5.8.2, but specific to red-on-a-line matchings. We will use it to additionally avoid the $\mathbf{H} \to \mathbf{T}$ transition. To state Lemma 5.8.3, we define the *upper cone* of two segments r_3b_3, r_4b_3 as the locus of the points that are separated from the horizontal line r_3r_4 by the two lines r_3b_3 and r_4b_3 (Figure 5.18(a)). We also define the *upper ray* of a segment as the open ray with the blue point as its origin, the segment as its direction, and going upwards (Figure 5.18(b)).

(a) (b)

Figure 5.18: (a) The upper cone of r_3b_3 and r_4b_3 is shaded. (b) The upper ray of the segment is dotted and highlighted.

Lemma 5.8.3. *In the* `Redonaline` *version, consider a set of three red-blue segments* $\{s, s_1, s_2\}$*, and the flip removing the pair of segments* s_1, s_2 *and inserting the pair of segments* s'_1, s'_2*.*

If the blue endpoint b *of* s *is not in any of the two upper cones of* s_1, s'_2 *and* s_2, s'_1*, then there always exists a tracking choice that avoids* $\mathbf{H} \to \mathbf{T}$ *for the pairs* s, s_1 *and* s, s_2 *while still avoiding* $\mathbf{H} \to \mathbf{X}$*.*

Proof. First, we check that there are only two possible upper cones defined by two segments of s_1, s_2, s'_1, s'_2. Indeed, only two pairs among them have a common blue point.

Then, we note that, for s, s_1 or s, s_2 to be in state $\mathbf{H}$, the red point r of s cannot be between r_1 and r_2, the red points of s_1 and s_2. Without loss of generality, we assume r to lie on the left side of r_1 and r_2.

For s, s'_1 or s, s'_2 to be in state $\mathbf{T}$, s has to cross at least one of the upper rays of s'_1 or s'_2.

The only two combinations of states for $\{s, s_1, s, s_2\}$ and $\{s, s'_1, s, s'_2\}$ which do not leave us the choice to avoid the $\mathbf{H} \to \mathbf{T}$ transition are $\{\mathbf{H}, \mathbf{T}\}$ and $\{\mathbf{T}, \mathbf{T}\}$, and $\{\mathbf{H}, \mathbf{X}\}$ and $\{\mathbf{T}, \mathbf{X}\}$. In any case, b must be in the right most of the two upper cones of segments s_1, s_2, s'_1, s'_2. More precisely, b lies in one of the three shaded regions of Figures 5.15 and 5.16. These three shaded regions also correspond to Figure 5.11 where case 1 is omitted but similar. The other cases are either not feasible geometrically, or with a possibility to make tracking choices so as to avoid transition $\mathbf{H} \to \mathbf{T}$. □

Proof of Theorem 5.8.1. We are now ready to prove Theorem 5.8.1.

Proof. Let $f(S)$ be the total number of flips performed by the algorithm on an n-segment input matching S and let $g(S)$ be the number of flips performed by the algorithm before the recursive calls. Let S_{rec} denote the matching before the recursive calls. The recursive calls take two submatchings of S_{rec} that we call S_1 and S_2, yielding the following recurrence relation.

$$f(S) = f(S_1) + f(S_2) + g(S)$$

Let $\bar{h}(S)$ be the number of $\mathbf{X}$-pairs plus the number of $\mathbf{T}$-pairs in a matching S, that is, the number of pairs that are not $\mathbf{H}$-pairs. Lemma 5.8.3 ensures that

$$g(S) \leq \bar{h}(S) - \bar{h}(S_{\text{rec}}) \leq \bar{h}(S) - \bar{h}(S_1) - \bar{h}(S_2).$$

Clearly, $f(\emptyset) = 0$. We suppose that, for all S' with less than n segments, we have $f(S') \leq \bar{h}(S')$. Then by induction we get

$$f(S) \leq \bar{h}(S_1) + \bar{h}(S_2) + \bar{h}(S) - \bar{h}(S_1) - \bar{h}(S_2) = \bar{h}(S).$$

Theorem 5.8.1 follows since $\bar{h}(S) \leq \binom{n}{2}$. □

State Tracking in the Convex Case. State tracking also applies to the widely studied convex case, providing a more conceptual proof of Theorem 3.2.2 [17].

Alternative proof of Theorem 3.2.2 [17]. In the convex case, the $\mathbf{T}$-state does not exist. Lemma 5.8.2 thus ensures that the number of $\mathbf{H}$-pairs increases by at least 1 at each flip. □

Chapter 6

Untangling with Insertion Choice

In this chapter, we devise strategies for insertion choice to untangle multisets of segments, therefore providing upper bounds on several versions of $\mathbf{d}^{\mathrm{I}}$. 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 `Separated` 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).

6.1 Upper Bound for Convex Position

Let $P = C = \{p_1, \ldots, p_{|C|}\}$ be a set of points in convex position sorted in counterclockwise order along the convex hull boundary (Figure 6.1(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.

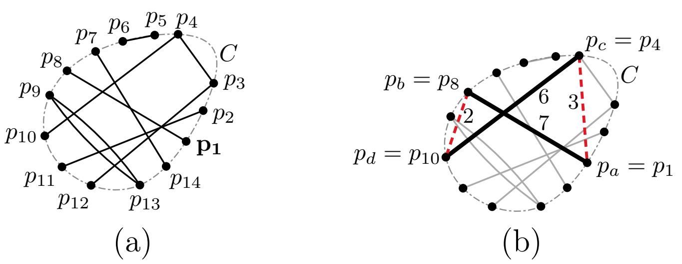

Figure 6.1: (a) A multigraph (C, S) with $|C| = 14$ points in convex position and $n = 9$ segments. (b) Insertion choice for Case 1 and 2 of the proof of Theorem 6.1.1. (c) Insertion choice for Case 3.

Theorem 6.1.1 ([31]). *Consider a multiset S of n segments with endpoints $P = C$ in convex position.*

In the `Convex Multigraph` *version, there exists an insertion strategy such that for any removal strategy, the resulting untangle sequence of S is of length $O(n \log |C|) = O(n \log n)$.*

[1]This definition resembles but is not exactly the same as the depth used in [17]. It is also similar to the crossing depth defined in Chapter 3.

In other words, we have the following upper bound:

$$\mathbf{d}^{\mathtt{I}}_{\mathtt{Convex\ Multigraph}}(n) = O(n \log |P|) = O(n \log n).$$

Proof. Let the potential function

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

As $\delta(s) \in \{1, \ldots, |C| - 1\}$, 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 6.1(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 6.1(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 $f' \leq 3f/4$. □

6.2 Upper Bound for Points Separated by Two Parallel Lines

In this section, we prove the following theorem, which is a generalization of Theorem 6.1.1. We extend our standard general position assumptions to also exclude pairs of endpoints with the same y-coordinate.

Theorem 6.2.1 ([31]). *Consider a multiset S of n segments with endpoints P 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. We define the parameter t as the sum of the degrees of the points in $T = T_1 \cup T_2$.*

In the `Separated Multigraph` *version, there exists an insertion strategy such that for any removal strategy, the resulting untangle sequence of S is of length $O(t\,|P| \log |C| + n \log |C|) = O(tn \log n)$.*

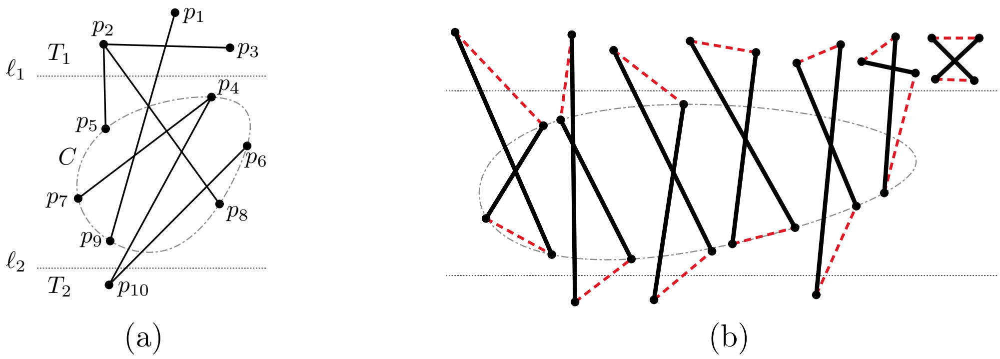

Figure 6.2: (a) Statement of Theorem 6.2.1. (b) Some insertion choices in the proof of Theorem 6.2.1.

In other words, we have the following upper bound:

$$\mathbf{d}^{\mathtt{I}}_{\mathtt{Separated\ Multigraph}}(n,t) = O(t\,|P|\log|C| + n\log|C|) = O(tn\log n).$$

Proof. We start by describing the insertion choice for flips involving at least one point in T. Let $p_1,\ldots,p_{|P|}$ 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_ap_b and p_cp_d. See Figure 6.2(b). As in the proof of Theorem 3.1.4 [19], we define the potential η of a segment p_ip_j as

$$\eta(p_ip_j) = |i-j|\,.$$

Notice that η is an integer between 1 and $|P|-1$. We define $\eta_T(S)$ as the sum of $\eta(p_ip_j)$ for $p_ip_j \in S$ 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 6.1.1. The potential function

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

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

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

Chapter 7

Untangling with Both Choices

In this chapter, we devise strategies for insertion and removal choices to untangle a multiset of segments, therefore providing upper bounds on several versions of $\mathbf{d}^{\mathtt{RI}}_{\mathtt{Allout}}$. 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 chapter, we work in the `Allout` version, i.e., with the following assumptions. We assume that the point set P is partitioned into $P = C \cup T$ where C is in convex position, and we define the parameter t as the sum of the degrees of the points in T. We additionally assume that the points of 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.

7.1 Upper Bound for Points Separated by Two Parallel Lines

In this section, we prove an upper bound on $\mathbf{d}^{\mathtt{RI}}_{\mathtt{Separated\ Multigraph}}$. Recall that, in the `Separated` version, 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 Theorem 3.2.13 [17, 31] and the $O(t\,|P|)$ bound from Theorem 3.1.4 [19] for t arbitrary segments. We extend our standard general position assumptions to also exclude pairs of endpoints with the same y-coordinate.

Theorem 7.1.1 ([31])**.** *Consider a multiset S of n segments with endpoints P 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. We define the parameter t as the sum of the degrees of the points in $T = T_1 \cup T_2$.*

In the `Separated Multigraph` *version, there exists an untangle sequence of S of length $O(n + t\,|P|) = O(tn)$.*

In other words, we have the following upper bound:

$$\mathbf{d}^{\mathtt{RI}}_{\mathtt{Separated\ Multigraph}}(n, t) = O(n + t\,|P|) = O(tn).$$

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 $p_1, \ldots, p_{|P|}$ 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 6.1.1 [31] and of Theorem 6.2.1 [31], we define the potential δ of a segment $p_i p_j$ as $\delta(p_i p_j) = |i - j|$.[1] 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 an algorithm to untangle a convex set with removal and insertion choice from Theorem 3.2.13, which performs $O(n)$ flips. $\square$

7.2 Upper Bound for Matchings with Points Outside a Convex

In this section, we prove an upper bound on $\mathbf{d}^{\mathtt{RI}}_{\mathtt{Allout\ Matching}}$. This upper bound may not hold for the `Multigraph` version as the proof uses a lemma specific the `Matching` version.

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 7.2.1. *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 7.1).

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 $p_1 p_2$ (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'.

[1] This definition resembles but is not the same as the depth used in [17].

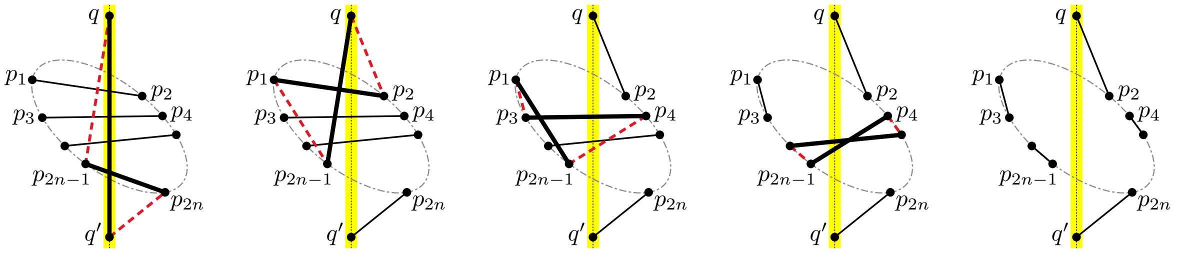

Figure 7.1: An untangle sequence of the subroutine to liberate the line qq' (with $n = 4$).

First flip. Lemma 7.2.2 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. □

Auxiliary Lemma of Section 7.2. In this section, we prove Lemma 7.2.2 used in the proof of Lemma 7.2.1.

Recall that, in the proof of Lemma 7.2.1, 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 7.2.2 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 7.2.2, we detail how to apply it to this context.

Lemma 7.2.2 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 7.2.2.

Lemma 7.2.2. *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. □

Proof of the Upper Bound

We are now ready to prove the following theorem, which only applies to matchings because it uses Lemma 7.2.1.

Theorem 7.2.3 ([31])**.** *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) forms a matching. We define the parameter t as the sum of the degrees of the points in T.*

In the `Allout Matching` *version, there exists an untangle sequence of length $O(t^3 n)$. In other words, we have the following upper bound:*

$$\mathbf{d}^{\texttt{RI}}_{\texttt{Allout Matching}}(n, t) = O(t^3 n).$$

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. As in the proof of Theorem 3.1.3, 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 7.2.1 to liberate s from every CC-segment using $O(n)$ flips. Let ℓ be the line containing s. Since Λ_ℓ cannot increase (Lemma 3.1.1), $\Lambda_\ell < t$ after Lemma 7.2.1, and there are $O(t^2)$ different TT_{in}-segments, it follows that Lemma 7.2.1 is applied $O(t^2)$ times, performing a total $O(t^2 n)$ 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, 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 7.2.4 and Lemma 7.2.5). 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 7.2.1).

There are two situation in which $\rho(S)$ may increase. One is when Lemma 7.2.1 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 $O(t^3\,|C|) = O(t^3 n)$.

CC-segments. By removal choice, we choose to flip the pairs of CC-segments last (except for the ones flipped in Lemma 7.2.1). As T is outside the convex hull of C, flipping two CC-segments does not create crossings with other segments (by the splitting lemma, i.e., Lemma 2.6.1 [19, 31]). Hence, we apply Theorem 3.2.13 to untangle the remaining segments using $O(n)$ flips. □

Auxiliary Lemmas of Section 7.2 In this section, we prove Lemma 7.2.5 and Lemma 7.2.4 used in the proof of Theorem 7.2.3.

Recall that, in the proof of Theorem 7.2.3, 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 7.2.4 and Lemma 7.2.5 to show that such a line ℓ always exist.

Indeed, by Lemma 7.2.4, 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 7.2.5 shows.

Next, we state prove Lemma 7.2.4 and Lemma 7.2.5.

Lemma 7.2.4. *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 7.2.5. *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. (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.)*

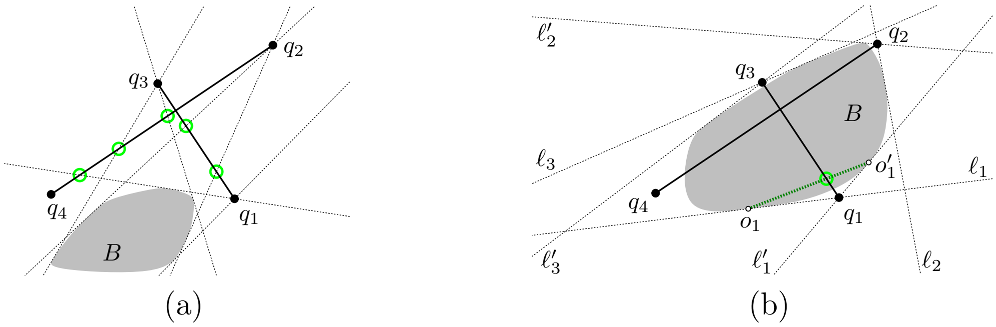

Figure 7.2: (a) In the statement of Lemma 7.2.5, we assert the existence of points, circled in the figure, which are the intersection of a line tangent to B and containing one of the points q_1, q_2, q_3. (b) In the proof of Lemma 7.2.5 by contraposition, we exhibit a point, circled in the figure, showing 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. 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 [49] 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, concluding the proof by contraposition. □

Chapter 8

Intractability of the Shortest Bipartite Untangle Sequence

This chapter is dedicated to the proof of the intractability of the shortest untangle sequence in the `Bipartite Matching` version (and of any constant factor approximation of it). Specifically, we prove the $\mathcal{NP}$-hardness of the following problem. Let $d(S)$ denote the length of the shortest untangle sequences of S.

Problem 1. *Let $\alpha \geq 1$ be a constant.*
Input: S, a set of segments with rational coordinates forming a bipartite matching.
Output: An untangle sequence starting at S of length at most α times $d(S)$.

We have the following theorem.

Theorem 8.0.1. *Problem 1 is $\mathcal{NP}$-hard for all $\alpha \geq 1$.*

8.1 Reduction Strategy

De Berg and Khosravi [37] showed that the rectilinear planar monotone 3-SAT problem (*RPM* 3-*SAT*) is $\mathcal{NP}$-hard. The RPM 3-SAT problem is a special case of the classic 3-SAT problem in which the clauses consist only of either all positive or all negative literals and the layout is planar (Figure 8.1). We reduce RPM 3-SAT to Problem 1. The key elements of the reduction are described next.

Given a planar embedding of an RPM 3-CNF formula φ (Figure 8.1), we construct a matching S_φ of polynomial size. The property of this matching S_φ is that its shortest untangle sequence has a length below a certain constant if φ is satisfiable and above α times this constant otherwise. Figure 8.2 shows the matching S_φ corresponding to the formula $\varphi = (x_1 \vee x_2 \vee x_3) \wedge (x_3 \vee x_4 \vee x_5) \wedge (x_3 \vee x_5 \vee x_6) \wedge (\overline{x_2} \vee \overline{x_3} \vee \overline{x_4})$ from Figure 8.1.

The aforesaid matching S_φ is built using two types of gadgets. The variable rectangles are replaced by variable gadgets (Figure 8.3). The clause rectangles together with the corresponding edges are replaced with padded clause gadgets. A padded clause gadget is represented in Figure 8.7 with plain segments. Throughout all the figures in this section, the dashed segments represent all the possibly created segments after any sequence of flips.

A *variable gadget* is a three-segment matching with two crossings. It allows for two possible flips, either of which produces a crossing-free matching, as shown in Figure 8.3. The flip generating the topmost segment stands for *false* ($x = 0$ in Figure 8.3), while the flip generating the bottom segment stands for *true* ($x = 1$).

A *clause gadget* is an OR gate with three inputs (Figure 8.8). The RPM 3-CNF clauses are either positive or negative. We describe the gadget for a positive clause, but the gadget for a negative clause can be defined analogously (by a vertical reflection). Three variable gadgets are the inputs of a clause gadget. In the crossing-free matching obtained for the clause gadget, the presence of the topmost segment (r_4b_7 in Figures 8.5, 8.6, 8.7, and 8.8) stands for a false output.

A *padding gadget* is a gadget that serves to force an arbitrarily large number k of flips if a clause is false. It consists of a series of k non-crossing segments (the plain segments in Figure 8.6, r_4b_7 aside). A *padded clause gadget* is a clause gadget coupled with a padding gadget in such a way that the presence of the output segment triggers k extra flips (Figure 8.7).

Let c be the number of clauses and v be the number of variables of the formula φ. If φ is *satisfiable*, then the shortest untangle sequence of S_φ has at most 5 flips per clause plus 1 flip per variable. In this case, we have $d(S_\varphi) \leq 5c + v$. We choose the size of the padding gadget so that a non-satisfied clause triggers $k = \alpha(5c + v) + 1$ flips. If the formula φ is *not satisfiable*, then at least one of the padding gadgets is triggered and $d(S_\varphi) > \alpha(5c + v)$.

8.2 The Problem to Be Reduced

In *RPM* 3-*SAT*, the *graph* of a conjunctive normal form (CNF) formula is the bipartite graph with the variables and clauses as vertices, and where there is an edge between a variable and a clause if and only if the clause contains the variable. A clause is said to be *positive* if it contains only positive variables; it is said to be *negative* if it contains only negative variables. A CNF formula is *monotone* if each clause is either positive or negative.

A *rectilinear planar monotone* 3-CNF (*RPM 3-CNF*) formula is a monotone formula with 3-variables per clause whose graph can be drawn with the following conventions (Figure 8.1). (i) The variables and the clauses are represented by axis-parallel non-overlapping closed rectangles. (ii) The variable rectangle centroids lie on the x-axis. (iii) The positive clause rectangles are above the x-axis, the negative ones, below. (iv) The edges connecting a variable to a clause are vertical line segments and do not cross any other rectangle. We call such a drawing a *planar embedding* of φ.

8.3 Variable Gadgets

A *variable gadget* is a three-segment matching built on the four endpoints of an axis-parallel rectangle as follows (Figure 8.3). The two leftmost endpoints of the rectangle are colored red, the two rightmost ones are colored blue. One of the segments of the matching is the diagonal joining the bottom left red point to the top right blue point. We add one red point on the vertical line cutting the rectangle in two symmetric halves,

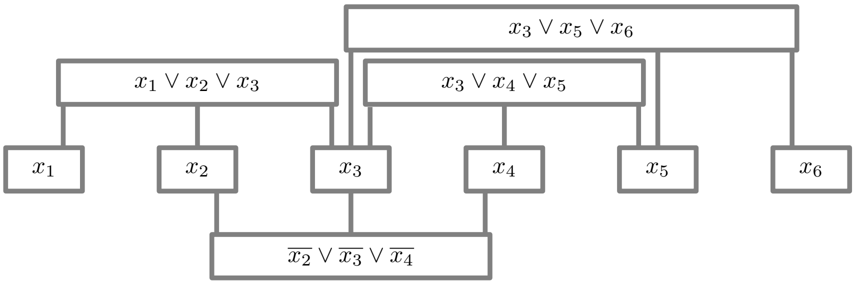

Figure 8.1: A planar embedding of an RPM 3-CNF formula φ.

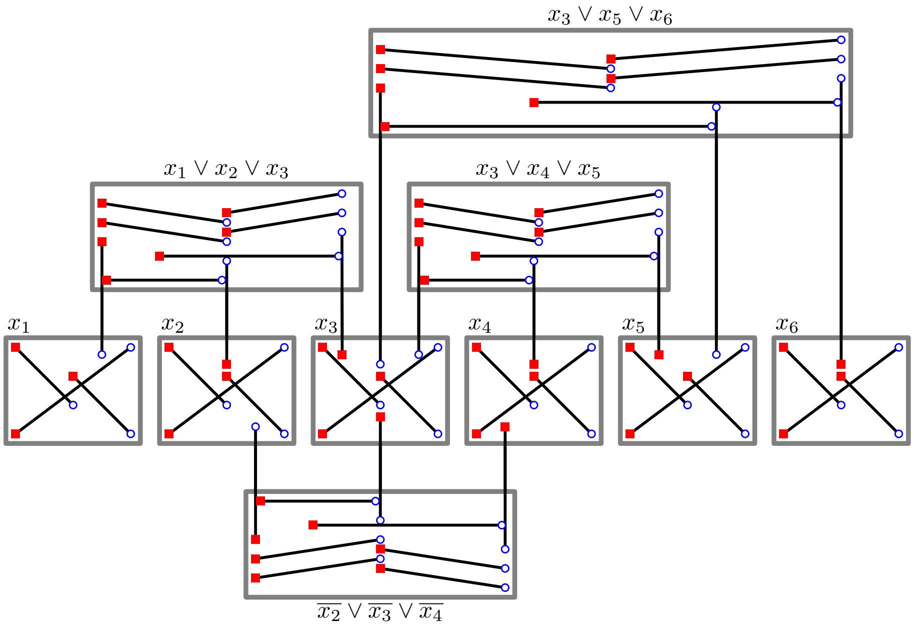

Figure 8.2: The matching S_φ of the formula φ from Figure 8.1.

just above the diagonal, in the inside of the rectangle. This red point is connected to the bottom right blue point. Similarly, we add one blue point on the same vertical, just below the diagonal. This blue point is connected to the top left red point.

We will refer to the triangle consisting of the three topmost points of a variable gadget as the *top triangle* of the variable gadget.

Lemma 8.3.1. *A variable gadget is the starting matching of exactly two untangle sequences of length* 1*, each untangle sequence ending in a distinct matching.*

Proof. It is straightforward to check the two possible cases. □

We can therefore represent each variable x of a propositional formula by a variable gadget. Assigning x to a truth value amounts to choosing one of the two possible untangle sequences, with the convention that the lower edge of the rectangle is present in the final matching if $x = 1$ (i.e., x is "true"), and that the upper edge of the rectangle is present if $x = 0$ (Figure 8.3).

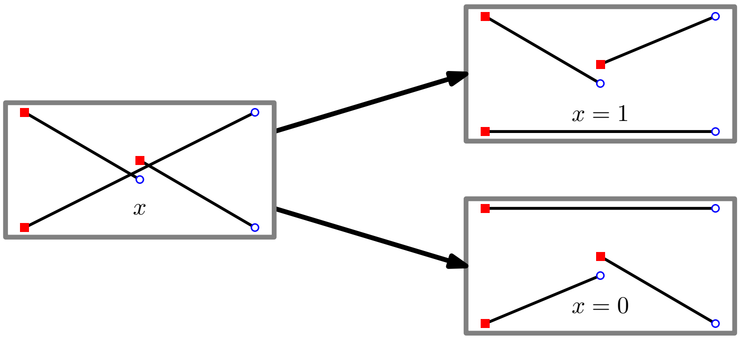

Figure 8.3: A variable gadget and its two untangle sequences.

8.4 OR Gadgets

An *OR gadget* consists of four three-segment matchings built on a common point set, say $\{r_1, r_2, r_2', r_3\}$ for the red points, and $\{b_1, b_1', b_2, b_3\}$ for the blue points, as follows (see the first matching in each of Figures 8.4(a), 8.4(b), 8.4(c), and 8.4(d), ignoring the dashed segments). The $0 \vee 0$ *matching* consists of the segments $r_1b_1', r_2'b_2, r_3b_3$, and only the first two are not crossing. The $0 \vee 1$ *matching* consists of the segments r_1b_1', r_3b_3, r_2b_2, and only the first two are crossing. The $1 \vee 0$ *matching* consists of the segments $r_2'b_2, r_3b_3, r_1b_1$, and only the first two are crossing. The $1 \vee 1$ *matching* consists of the segments r_1b_1, r_2b_2, r_3b_3, and is crossing-free. In addition to these constraints, the point set also satisfies the following ones. The following three matchings are crossing-free: $\{r_1b_2, r_2'b_3, r_3b_1'\}$, $\{r_1b_3, r_2b_2, r_3b_1'\}$, and $\{r_1b_1, r_2'b_3, r_3b_2\}$. In each of the following two matchings, only the first two segments are crossing: $\{r_1b_3, r_2'b_2, r_3b_1'\}$, and $\{r_1b_1', r_3b_2, r_2'b_3\}$.

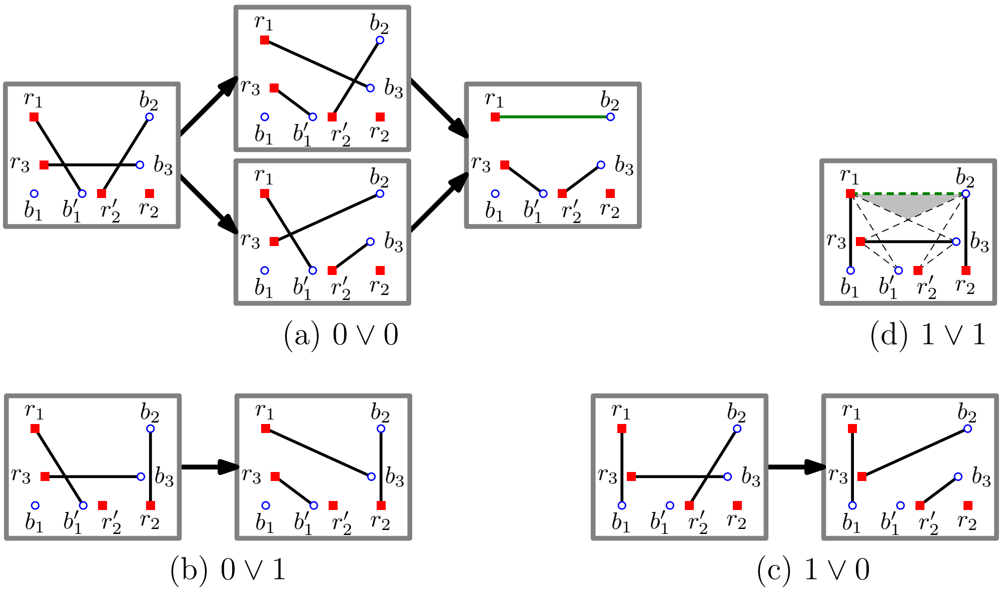

Figure 8.4: The four matchings of an OR gadget, with their untangle sequences.

Note that, in any of the four matchings of an OR gadget, there is one unused blue point and one unused red point. If the unused blue point is b_1 (respectively b_1'), we say that the *left input* of the OR gadget is 0 (respectively 1). Similarly, if the unused red point is r_2 (respectively r_2'), we say that the *right input* of the OR gadget is 0 (respectively 1). To complete the similarity with a logical gate, we also define the *output* of the OR gadget as 0 if the segment r_1b_2 is present in all the final matchings of any untangle sequence starting at the OR gadget and as 1 if the segment r_1b_2 is absent of all the same final matchings. The output is undefined otherwise. The following lemma states that the truth table of the logical gate associated with an OR gadget is indeed the one of an OR gate.

We will refer to the smallest of the triangles consisting of the segment r_1b_2 and induced by all the other segments we have mentioned in the definition of an OR gadget as the *top triangle* of the OR gadget. It is the shaded triangle in Figure 8.4(d).

Lemma 8.4.1. *The output of an OR gadget is always well defined, and is* 0 *if and only if the two inputs of the OR gadget are both* 0*. More precisely, we have the following.*

1. *The* $0 \vee 0$ *matching is the starting matching of exactly two untangle sequences, each of length* 2*, and ending at the same matching containing the upper segment* r_1b_2 *(Figure 8.4(a)).*

2. *The* $0 \vee 1$ *matching is the starting matching of a unique untangle sequence of length* 1 *ending at a matching excluding the upper segment* r_1b_2 *(Figure 8.4(b)).*

3. *The* $1 \vee 0$ *matching is the starting matching of a unique untangle sequence of length* 1 *ending at a matching excluding the upper segment* r_1b_2 *(Figure 8.4(c)).*

4. *The* $1 \vee 1$ *matching is already crossing free. It excludes the upper segment* r_1b_2 *(Figure 8.4(d)).*

Proof. For each of the four $x \vee y$ matchings with $x, y \in \{0, 1\}$, we enumerate all the possible untangling sequences. These sequences are all shown in Figure 8.4. Lemma 8.4.1 then follows. □

8.5 Clause Gadgets

A *clause gadget* consists of two OR gadgets, the output of the first one being "connected" to the left input of the second one (Figure 8.5). More precisely, a clause gadget is built on seven red points, say $r_4, r_5, r_6, r_7, r_8, r_{10}, r_{11}$, and six blue points, say $b_4, b_5, b_6, b_7, b_8, b_9$ such that the following maps correspond to two OR gadgets (using the OR gadget previous notations), and such that r_8 lie in the inside of the top triangle of the first OR gadget and is the only overlap between the two OR gadgets.

$$\text{First OR gadget: } (r_4, b_4, r_5, b_5, r_6, b_6, b_9, r_{10}) \mapsto (r_1, b_1, r_2, b_2, r_3, b_3, b_1', r_2').$$
$$\text{Second OR gadget: } (r_4, b_6, r_7, b_7, r_8, b_8, b_5, r_{11}) \mapsto (r_1, b_1, r_2, b_2, r_3, b_3, b_1', r_2'),$$

with the exception that the segment r_6b_5 may also play the role of r_1b_1.

Similarly to an OR gadget, a clause gadget consists of 2^3 matchings, namely the $x \vee y \vee z$ *matchings* with $x, y, z \in \{0, 1\}$. We define the *left*, *middle*, and *right input* of a clause gadget as the left input of the first OR gadget, the right input of the first OR gadget, and the right input of the second OR gadget. We define the *output* of a clause gadget as the output of the second OR gadget.

Note that the middle input segment, i.e., the vertical segment lying in between the two other vertical segments (r_5b_5 in Figure 8.5), need not be evenly placed between the left input segment and the right input segment. This feature is used to build clause gadgets with non-consecutive variables, such as the topmost clause gadget in Figure 8.2.

The idea is to have a $0 \vee 0 \vee 0$ clause gadget in S_φ for each clause in φ. As we will see next, in the beginning of an untangling sequence starting at S_φ, each input may be set to 1 or may be kept as 0, changing the $0 \vee 0 \vee 0$ clause gadget into one of the $x \vee y \vee z$ matchings with $x, y, z \in \{0, 1\}$.

The following lemma states that the truth table of the logical gate associated with a clause gadget is indeed the expected one.

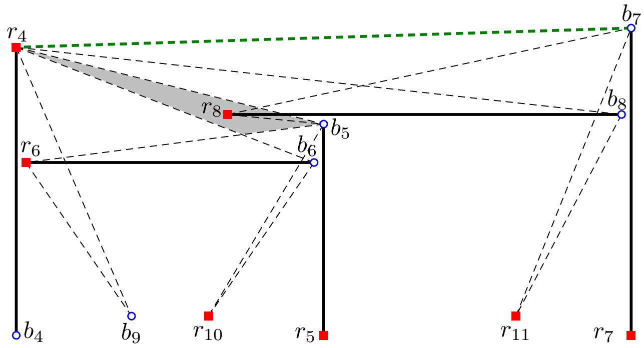

Figure 8.5: A clause gadget. The $0 \vee 0 \vee 0$ matching is drawn with plain segments.

Lemma 8.5.1. *The output of a clause gadget is always well defined, and is* 0 *if and only if the three inputs of the clause gadget are all* 0*. More precisely, we have the following.*

1. *All the untangle sequences starting at the* $0 \vee 0 \vee 0$ *matching are of length* 4*, and they end at the same matching containing the upper segment* r_4b_7*.*

2. *All the untangle sequences starting at each of the* $x \vee y \vee z$ *matchings, where exactly one of* $x, y,$ *or* z *is* 1*, are of length* 2*, and they end at matchings excluding the upper segment* r_4b_7*.*

3. *The unique untangle sequence starting at each of the* $x \vee y \vee z$ *matchings, where exactly two of* $x, y,$ *and* z *are* 1*, is of length* 1*, and it ends at a matching excluding the upper segment* r_4b_7*.*

4. *The* $1 \vee 1 \vee 1$ *matching is already crossing free, and it excludes the upper segment* r_4b_7*.*

Proof. It is a consequence of Lemmas 8.4 and of the fact that the OR gadgets are connected so as to not interfere. Indeed, by construction, r_8 lies in the inside of the top triangle of the first OR gadget, and is the only overlap between the two OR gadgets. This ensures that all untangle sequences never give rise to an extra crossing that does not already belong to one of the two OR gadgets. In Figure 8.5, we have drawn with dashed line segments all the possible created segments during any possible untangle sequence. □

8.6 Padding Gadgets

Let k be a non-negative integer. A *k-padding gadget* triggered by the segment s consists of two matchings built by induction as follows.

The first matching, denoted S_k, contains s ($s = r_4b_7$ in Figure 8.6) and is called the *triggered matching* of the padding gadget (s creates a crossing). The second matching is called the *non-triggered matching*, and is deduced from the triggered one by removing s (it is crossing free).

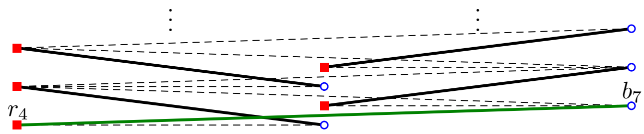

Figure 8.6: The triggered matching of a padding gadget.

If $k = 0$, then the triggered matching of a k-padding gadget consists of only the segment s. If $k \geq 1$, then the triggered matching of a k-padding gadget consists of S_{k-1}, the triggered matching of a $(k-1)$-padding gadget, to which we add one new segment crossing only the last created segment of the only untangle sequence starting at S_{k-1} (Figure 8.6, the dashed segments are all the possible created segments in the unique untangle sequence).

Lemma 8.6.1. *Let k be a non-negative integer. There exists a unique untangle sequence starting at the triggered matching of a padding gadget, and its length is k. The non-triggered matching of a padding gadget is already crossing free.*

Proof. The definition of a k-padding gadget yields Lemma 8.6.1. □

We complete each clause gadget with a padding gadget in order to penalize a non-satisfied clause by an arbitrary long untangle sequence (Figure 8.7). Notice that a padded clause gadget can be arbitrarily scaled and that the position of a clause rectangle is only constrained by the planar embedding of φ.

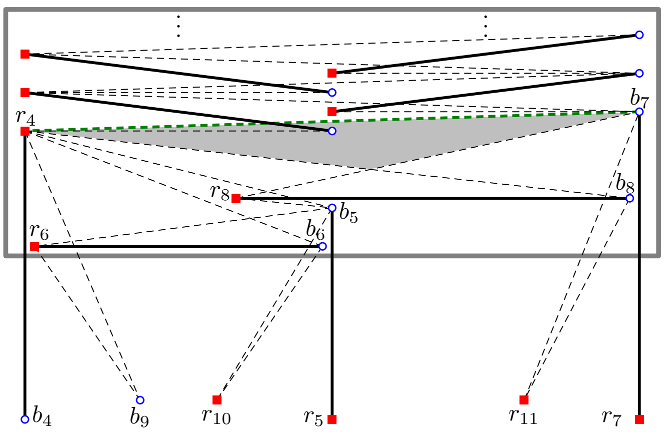

Figure 8.7: A clause gadget connected to a padding gadget.

8.7 Matching Computation

We now describe, given a planar embedding of φ, the construction steps of the matching S_φ. Without loss of generality, we only specify the construction of the positive clauses, the construction of the negative clauses being similar.

We need the following definitions for the description. The *top vertical points* of a positive clause gadget are the topmost endpoints of the vertical segments (e.g. $b_{14}, r_4, r_{12}, b_5, b_7$ in Figure 8.8). Similarly, the *bottom vertical points* are the bottom endpoints of the same vertical segments (e.g. $r_{14}, b_4, b_{12}, r_5, r_7$ in Figure 8.8). Let $\overline{p}$ be a top vertical point. The *horizontal segment of* $\overline{p}$ is the horizontal segment lying below $\overline{p}$ which is the closest to $\overline{p}$. Finally, we define the *substitute point* of $\overline{p}$ as the endpoint of the horizontal segment of $\overline{p}$ which is the closest to $\overline{p}$ (e.g. r_6 is the substitute point of r_4 in Figure 8.8).

The construction steps of the matching S_φ are the following.

1. Place a clause gadget connected to a k-padding gadget in each clause rectangle, and a variable gadget in each variable rectangle, with appropriate scaling.

2. Connect each clause gadget to its corresponding three variable gadgets with the three vertical segments of the clause gadget aligned with the corresponding vertical edges of the planar embedding of φ.

3. Adjust the x-coordinates of the vertical segments of each variable gadget to have the top vertical points and the two topmost points of the variable gadget, all in convex position (e.g. in Figure 8.8, r_{12} is on the right of the segment $r_4 b_9$).

4. Adjust the y-coordinates of the bottom vertical points in the top triangle of each variable gadget so as to place them and the two topmost points of the variable gadget in convex position.

5. Let $\overline{p}$ be a top vertical point which is not the highest of a variable gadget (e.g. $\overline{p} = r_{12}$ in Figure 8.8). Let $\underline{p}$ be the corresponding bottom vertical point (e.g. $\underline{p} = b_{12}$). Let $\overline{q}$ be the top vertical point immediately above $\overline{p}$ (e.g. $\overline{q} = r_4$). Let $\underline{q}'$ be the point immediately above $\underline{p}$, taken among the bottom vertical points together with the two topmost points of the variable gadget (e.g. $\underline{q} = b_9$). Adjust the x-coordinate of $\tilde{p}$, the substitute point of $\overline{p}$ (e.g. $\tilde{p} = r_{13}$), so that $\tilde{p}$ lies in the triangle $\overline{p}\underline{q}\overline{q}$ (e.g. a shaded triangle in Figure 8.8; segment $r_{13}b_{13}$ must not cross $r_4 b_9$, but it has to cross $r_{12} b_9$).

We have the following lemma.

Lemma 8.7.1. *Let φ be an instance of RPM 3-SAT with c clauses and v variables. Let k be a non-negative integer, polynomial in c and v. The matching S_φ with k-padding gadgets is computed in polynomial time in c and v.*

Proof. The number of operations in any execution of these construction steps is linear in c and v. The coordinates of the points of S_φ are rational numbers with $O(\log n)$ bits. □

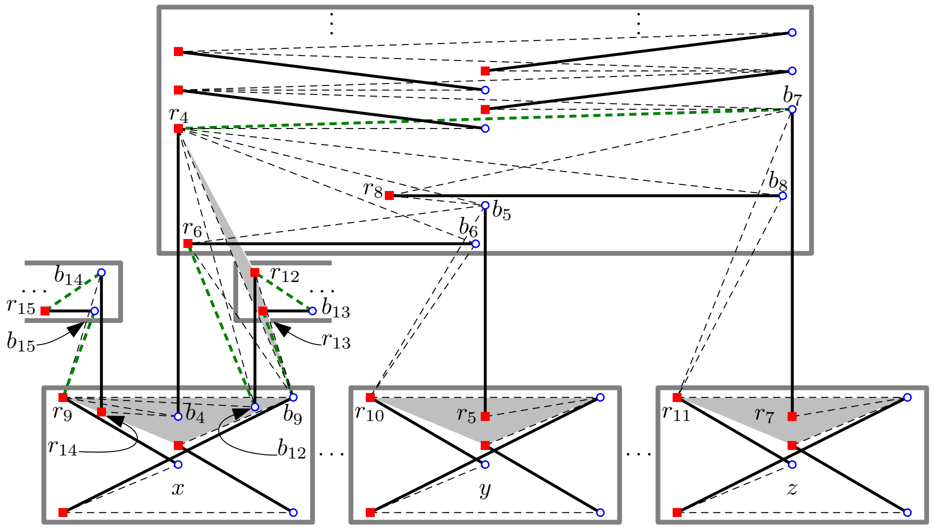

Figure 8.8: A padded clause gadget connected to x, y, z, with branching on x.

8.8 Branching

The following lemma ensures that the connection of multiple vertical segments to a same variable gadget always triggers all the corresponding clause gadgets. We start with some definitions.

The set consisting of the top segment of a variable gadget set to true, together with the vertical segments crossing it, and their horizontal segment is called a *branching matching* (such as drawn in Figure 8.9(b) with plain segments). The bottom vertical points of a branching matching, listed from left to right, always consist of a certain number, say a, of red points followed by a certain number, say b, of blue points. We say that such a branching matching has parameters a, b. These matchings have the following property.

Lemma 8.8.1. *All the untangle sequences starting at a branching matching with parameters a, b have length $2(a + b)$ and end at the same crossing-free matching (e.g. the segments $r_{15}b_{14}, r_9b_{15}, r_{14}b_4, r_4b_6, r_6b_{12}, r_{12}b_{13}, r_{13}b_9$ in Figure 8.9(b)).*

Proof. First note that a simplified version of this result has been proven in Theorem 3.2.12 [19]. This simplified version amounts to forget all the horizontal segments, except the top segment of the variable gadget (Figure 8.9(c)).

It is useful to start by proving this simplified version before Lemma 8.8.1. We do an induction on $a + b$. The base case is trivial, but it provides the possible positions of created segments in the untangle sequences (the dashed segments in Figure 8.9(c)). The inductive case relies on the fact that the points are in convex position. Indeed, after any flip, the two created segments play the role of the initial horizontal segment because convex position ensures that any of the non-flipped vertical segments will cross exactly one of the two created segments, and that no extra crossing is created. The induction hypothesis then applies on both right and left submatching whose convex hulls are now disjoint.

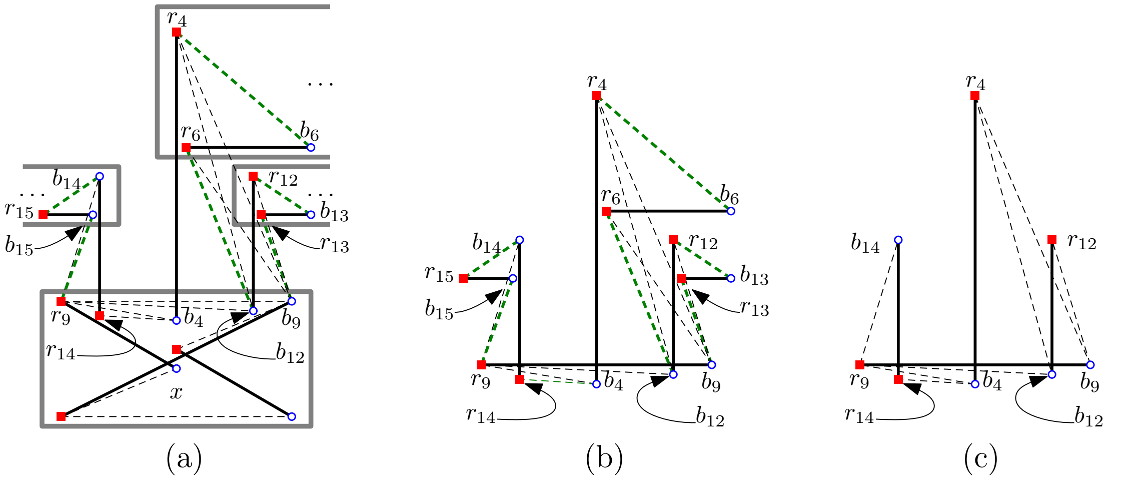

Figure 8.9: Three views of a branching matching.

We now address the issue where each vertical segment is paired with its horizontal segment. Recall that at step 5 of the construction of S_φ, we have adjusted the x-coordinate of $\tilde{p}$, the substitute point of each top vertical point $\overline{p}$, so that $\tilde{p}$ lies in the triangle $\overline{p}q\overline{q}$. This ensures that each substitute point can play the role of its corresponding top vertical point from whenever the corresponding horizontal segment has been flipped in an untangle sequence. □

8.9 Result

The RPM 3-SAT instance being encoded in the matching S_φ, we have the property that the shortest untangling sequence of S_φ is short if the instance φ is satisfiable, and long otherwise.

Lemma 8.9.1. *We have the following case distinction.*

- *φ is satisfiable if and only if there exist untangle sequences starting at S_φ which do not trigger any padding gadget, in which case $d(S_\varphi)$ is at most $v + 5c$.*

- *φ is not satisfiable if and only if all untangle sequences starting at S_φ trigger at least one padding gadget, in which case $d(S_\varphi)$ is at least $v + 7 + k$ where k is arbitrarily large.*

Proof. It is consequence of Lemmas 8.3.1, 8.5.1, 8.6.1 and 8.8.1, as we examine the longest possible untangle sequences of S_φ which do not trigger any padding gadget, and the shortest possible untangle sequences of S_φ which trigger at least one padding gadget. In any case, v flips will be performed, one per variable (Lemma 8.3.1).

In the case where no padding gadget is triggered, the length of the longest possible untangle sequences starting at a clause gadget connected to three variable gadgets is 5, and is obtained by adding 3, the length of the untangle sequences of a $0 \vee 0 \vee 1$ matching, and 2, for the two connections to the negative variables. Counting 5 flips per clause yields $v + 5c$.

If at least one padding gadget is triggered, this very padding gadget generates k flips. In this case, the length of the shortest possible untangle sequences starting at a clause gadget connected to three variable gadgets and which is known to trigger its padding gadget is 7, and is obtained by adding 4, the length of the untangle sequences of a $0 \vee 0 \vee 0$ matching, and 3, for the three connections to the negative variables. All the other cause gadgets may be set to their $1 \vee 1 \vee 1$ matching, adding no flip to the shortest untangle sequence, the length of which is thus $v + 7 + k$. □

We now prove Theorem 8.0.1, reducing RPM 3-SAT to Problem 1. Let φ be an instance of RPM 3-SAT with c clauses and v variables. We build the matching S_φ, which serves as an instance of Problem 1, choosing $k = \alpha(v+5c)+1$. As k is polynomial in the size of the input (α is a constant), the computation of the matching S_φ is polynomial (Lemma 8.7.1).

By hypothesis, we compute an untangle sequence starting at S_φ of length ℓ at most $\alpha d(S_\varphi)$. We decide that φ is satisfiable if $\ell \leq \alpha(v + 5c)$, and that φ is not satisfiable if $\ell > \alpha(v + 5c)$.

Indeed, Lemma 8.9.1 ensures the following. If φ is satisfiable, then the length of the shortest untangle sequence of S_φ is at most $v + 5c$. Otherwise the length of the shortest sequence is at least $v + 7 + k \geq k = \alpha(v + 5c) + 1$. This ends the reduction, and proves Theorem 8.0.1.

Chapter 9

Conclusion

This chapter is less a conclusion to this dissertation than an introduction to possible future developments. For each choice (i.e., for no choice, removal choice, insertion choice, and both choices) and for intractability, we explain the difficulties to extend or improve known results sometimes providing counter-examples to some intuitive ideas, we present some open problems which might be stepping stones to tackle the main open problems, and we discuss about what we believe could be promising ideas.

9.1 Untangling with No Choice

When no choice is available, the only approach so far has been to find "good" potentials to measure how far a multiset of segments is from being crossing free. To specify what "good" means, we define (in the Π version for some property Π) the *faithful* potential of a multiset S of segments $\mathbf{d}_{\Pi}^{\emptyset}(S)$ as the length of longest untangle sequence of S. Obviously, $\mathbf{d}_{\Pi}^{\emptyset}(S)$ decreases by at least one at each flip, and any other potential does not provide a better upper bound than $\max_S \mathbf{d}_{\Pi}^{\emptyset}(S)$. The issue is then to prove a good upper bound on $\mathbf{d}_{\Pi}^{\emptyset}(S)$. Now, we see that a *good* potential is one that is as close to the faithful potential as possible while still admitting a good upper bound for which we have a proof. In the following, we first present some variations of the potentials used in the previous chapters with some of their interesting consequences.

The potential of one point. The potential of a point $p \in P$ is defined as the sum of the potentials of the lines through p, i.e., the sum of the $\Lambda_{pp'}$ for all $p' \in P \setminus \{p\}$ (Λ is defined just before Lemma 3.1.1 [83]). This potential shows that the point p is involved in at most n^2 flips in an untangle sequence. Using the fact that there is a zone defined by each flip f such that all the points both in P and in this zone define f-critical lines (see the proof of Theorem 4.3.1 [30] for the definition of an f-critical line, see Lemma 4.5.3 [30] for another use of this fact), we can show that a ray from p containing the segment adjacent to p and rotating in the direction minimizing its angle of rotation sweeps at most n^2 points of P in an untangle sequence.

The potential of two points. There are three alternatives to define the potential of two points.

The potential of two points p_1, p_2 may be defined as the potential of the line p_1p_2, i.e., $\Lambda_{p_1p_2}$. This potential decreases by at least 1 at each flip removing the segment p_1p_2 and takes at most n distinct values, thus proving that a given segment is removed at most $n-1$ times (and inserted at most n times).

Similarly, the potential of two points p_1, p_2 may be defined as the sum of the potentials of two lines crossing the segment p_1p_2 which are close enough to the line p_1p_2. This potential decreases by at least 2 at each flip removing the segment p_1p_2 and takes at most $2n-1$ distinct values, thus proving that a given segment is removed at most $n-1$ times (and inserted at most n times). This potential proved handy in the proof of Theorem 3.1.3 to cope with the absence of any general position assumption.

Alternatively, the potential of two points p_1, p_2 may be defined as the sum of the potentials of two lines being close enough to the line p_1p_2 while being parallel to p_1p_2. This potential decreases by at least 2 at each flip inserting the segment p_1p_2 and takes at most $2n-1$ distinct values, thus proving that a given segment is removed at most $n-1$ times (and removed at most n times). This potential is introduced in [19] to prove Theorem 3.1.3 in the `Matching` version (assuming general position).

The potential of one segment. In the `Convex` version, we may define the potential of one segment s of S as the number of segments of S crossing s. We have found no way to generalize this type of potential to the non-`Convex` versions.

The potential of a convex partition. Note that a line defines a convex partition of the plane. One idea is to try to define a potential of a convex partition which is not necessarily induced by a line. Our tries all led to functions increasing in some flip situations.

A potential based on the cyclic order of the points. The potential Φ based on the cyclic order of $P \setminus \{p\}$ around each point p used in the proof of Theorem 4.4.1 in the `Redonaline` version is very interesting. Indeed, the idea is to see untangling as a generalisation of sorting. Moreover, counting the number of points in P that "see" a given **T**-pair as an **X**-pair provides a "measure" of how much this **T**-pair is from being **X** (or **H**). For instance, this observation can be used to relax the requirement that the removed segments cross in the definition of a flip to the requirement that less points in P "see" the inserted pair has an **X**-pair than the removed pair. Unfortunately, the potential Φ does not generalise well to the other versions. A possible reason is that the point set order type (see Section 1.5 for the definitions of the order type and crossing type of a point set) is not determined by this list of cyclic orders (see [5] for a reconstruction of all the order types of a given list of cyclic orders). It seems also unlikely that the point set crossing type would be determined by this list of cyclic orders (see for example [73] for an algorithm reconstructing the crossing type of a point set from its planar spanning trees organised in a reconfiguration graph).

A potential based on state tracking. As we have seen in Section 5.8, the state tracking framework sheds a new light on the untangling problems. For instance, it is possible to see the line potential Λ_L (where L is the set of the lines defined by any pair of points in the point set P) as counting the transitions $\mathbf{X} \to \mathbf{H}$. In fact, Λ_L

also decreases by at least one for each $\mathbf{H} \to \mathbf{T}$ and each $\mathbf{T} \to \mathbf{X}$ (via critical lines). As we already know that each $\mathbf{X}$-pair may be flipped at least once, we may focus on counting how many times each pair is flipped beyond once. We may therefore hope that some lines in L are not needed anymore if we focus on counting $\mathbf{H} \to \mathbf{T}$ (or $\mathbf{T} \to \mathbf{X}$) transitions instead of $\mathbf{X} \to \mathbf{H}$. Unfortunately, all the lines are needed in general. In the `Bipartite` version, only the lines defined by to points of P of the same color are needed (instead of the lines defined by to points of P of different colors), but their number is still quadratic, leading to the same cubic upper bound again. In fact, in the context of bounds parameterized by t, the first proof of Theorem 4.3.1 was found in the case where the points of T are inside the convex hull of C using the state tracking approach to count $\mathbf{H} \to \mathbf{T}$ transitions with fewer lines than in L.

Searching head-on for good potentials is one strategy. In the following, we present interesting remarks which are less focused on potentials.

The crossing type of a flip sequence. Similarly to the crossing type of a point set, we define the *labeled crossing type* of a flip sequence as the set of the pairs of pairs of labels of points which corresponds to the set of pairs of segments which are removed at some point in the flip sequence. Two flip sequences have the same *crossing type* if there exists a relabeling of the points such that the two flip sequences have the same labeled crossing type. This definition is useful to realize that

- a lot of the flip sequences we may come up with have the same crossing type (hence at most a quadratic length) as a flip sequence with endpoints in convex position,
- when we are only interested in some multisets of segments on a given point set, the problem depends only on the union of the crossing types of all the flip sequences of the multisets of interest (which may be even less constrained than the crossing type of the point set itself).

The intuition that all the flip sequences have the same crossing type as some flip sequence with endpoints in convex position is disproved by Theorem 4.2.1.

Reductions. In Lemma 2.3.1, Lemma 2.3.2, Lemma 2.3.3, and Theorem 4.1.1, we showed a relationship among different versions of $\mathbf{d}^{\emptyset}$. The latter result shows how upper bounds in the `Matching` version can be easily transferred to different versions. But they can also be applied to transfer lower bounds among different versions. For example, the lower bound of $\frac{3}{2}\binom{n}{2} - \frac{n}{4}$ in the `Bipartite Matching` version from Theorem 4.2.1 [33] implies a lower bound of $\frac{1}{3}\binom{n}{2} - \frac{n}{2}$ in the `Cycle` version. It is not clear if the constants in the `Cycle` version lower bound may be improved, perhaps by a more direct approach (or perhaps the lower bounds are not even asymptotically tight). However, we showed that all these versions are related by constant factors.

Upper bound without multiplicity. The upper bound from Theorem 4.5.1 on the number of distinct flips may be useful if one of the following open problems is solved.

- What is the maximal multiplicity of a flip removing a given segment p_1p_2 in an untangle sequence? We know that it is at most $n-1$ thanks to the line potential $\Lambda_{p_1p_2}$. Any sub-linear upper bound would lead to a sub-cubic upper bound on $\mathbf{d}^{\emptyset}_{\texttt{Multigraph}}$.
- Is it possible to devise a removal strategy which guaranties that the multiplicity of each flip is sub-linear? A yes would improve the cubic upper bound on $\mathbf{d}^{\mathtt{R}}_{\texttt{Multigraph}}$.

The proof of Theorem 4.5.1 is based on the $O(n^2k^2)$ bound from Lemma 4.5.3. A better analysis of the dual arrangement could potentially improve this bound, perhaps to $O(n^2k)$.

An inductive approach. If we add a new segment to a crossing-free multiset of segments, what is the maximum length of an untangle sequence? Notice that a sub-quadratic bound would lead to a sub-cubic bound for $\mathbf{d}^{\emptyset}_{\texttt{Multigraph}}$.

A tight quadratic lower bound. All our small scale experiments have so far shown that the quadratic lower bound seems to be asymptotically tight. Specifically in the `Bipartite` version, all our small scale experiments have so far shown that the lower bound of Theorem 4.2.1 seems to be exactly tight for even n. Apparently, bipartite matchings with endpoints in general position do not lead to untangle sequences longer than in the `Redonaline` version.

9.2 Untangling with Removal Choice

Devising removal strategies is strongly related to the motivation of untangling TSP tours. In this section, we discuss about lower bounds, what to do with the results in the `Redonaline` version, how counting flips without multiplicity could become useful, a recursive approach, the irregular dependencies in t in parameterized bounds, and last but not least we present work in progress involving state tracking to prove a quadratic upper bound on $\mathbf{d}^{\mathtt{R}}_{\texttt{Bipartite Matching}}$.

A tight linear lower bound. All our small scale experiments have so far shown that the linear lower bound seem to be asymptotically tight. In the following, we forget about this remark to discuss about how to prove lower bounds.

First note that most of the removal strategies are not known to be optimal. In fact, only the `Convex` version algorithms are shown to have optimal worst case flip complexity, except in the `Multigraph` version. It is uncertain whether the $O(n \log n)$ upper bound or the linear lower bound on $\mathbf{d}^{\mathtt{R}}_{\texttt{Convex Multigraph}}$ can be improved. Besides, in the `Convex Bipartite` version, there are hints (explained in Section 9.5, see footnote[3]) for the lower bound of Theorem 5.1.1 to be exactly tight for even n.

Improving the lower bound amounts to devising a strategy in place of the adversary, using insertion choice to make the untangle sequence as long as possible. As some multisets of segments do not admit untangle sequence with super-linear length, the choice of the initial multiset of segments is important. Moreover, to fully benefit from this choice, we have to use insertion choice to decrease as little as possible a potential for which we know a high lower bound on its initial value and an low upper bound its the final value. The only three potentials known to have a low upper bound on crossing-free multisets of segments (in fact, they all equal 0 on crossing-free multisets) are the number of crossings χ and the sum, or the product, of the crossing depths $\delta_\times(s)$ of the segments s in the multiset (the potential involving the crossing depths are specific to the `Convex` version). There are flip situations where both insertion choices lead to a big decrease of any of these three potentials. For instance, the potential χ decreases by $n-2$ for any flip of a multiset having each of its segments crossing all its others (with endpoints in convex position). Avoiding these situations or proving that they are not frequent (for example, by a similar argument than in the proof of Theorem 6.1.1 [31]) would lead to a lower bound on $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex\ Multigraph}}$.

From the `Redonaline` version to generalizations. Generalizing the algorithm used in the proof of Theorem 3.3.1 to bipartite point sets with the red points below the x-axis and the blue points above using ray shooting ideas to generalize the path u did not work.[1]

Generalizing the algorithm used in the proof of Theorem 5.8.1 did not work either as Lemma 5.8.3 relies on the red-on-a-line setting. It may generalize to some deformation of this setting like red-on-a-convex for instance.

Upper bound without multiplicity. Recall that a removal strategy which guaranties that the multiplicity of each flip is sub-linear together with the upper bound from Theorem 4.5.1 on the number of distinct flips may improve the cubic upper bound on $\mathbf{d}^{\mathtt{R}}_{\mathtt{Multigraph}}$.

A recursive approach. A recursive removal strategy which untangles a new segment added to a crossing-free multiset of segments in sub-quadratic time would lead to a sub-cubic bound for $\mathbf{d}^{\mathtt{R}}_{\mathtt{Multigraph}}$.

The dependency on t. In the upper bound of Theorem 5.5.2 [31], the quadratic dependency on t provides a continuous transition between the upper bounds in the `Convex` versions (as in Theorem 5.2.1) and the cubic upper bound for general point sets from Theorem 3.1.3. In the upper bounds of Theorem 5.4.2 [31] and Theorem 5.6.1 [31], the linear dependency on t would provide a continuous transition between the upper bounds in the `Convex` versions (as in Theorem 5.2.1) and a conjectured quadratic upper bound for general point sets (see the next paragraph "State tracking to prove a quadratic upper bound on $\mathbf{d}^{\mathtt{R}}_{\mathtt{Bipartite\ Matching}}$."). In the upper bound of Theorem 5.7.1, the exponential dependency on t seems more to emerge for technical reasons than to reflect an intrinsic aspect of the problem.

State tracking to prove a quadratic upper bound on $\mathbf{d}^{\mathtt{R}}_{\mathtt{Bipartite\ Matching}}$. The author believe it is possible to use state tracking to prove a quadratic upper bound on $\mathbf{d}^{\mathtt{R}}_{\mathtt{Bipartite\ Matching}}$. Next, we provide some insight on this intuition.

[1]This was discussed at the workshop [86].

As we have seen in the proof of Theorem 3.2.1, there is a correspondence between the permutations of n elements and the bipartite matchings on a given bipartite point set where the red (respectively blue) points are labeled $r_1, \dots, r_n$ (respectively $\bar{r}_1, \dots, \bar{r}_n$). Specifically, a bipartite matching S defines a permutation σ where, for each segment $r_a\bar{r}_b$ in S, a is mapped to b. For such a bipartite matching, the flips correspond to a subset of the transpositions. Specifically, a flip (in the `Bipartite` version) removing the segments $r_a\bar{r}_b, r_c\bar{r}_d$ corresponds to the transposition (ac) (i.e. the transposition mapping a to c and vice versa). A flip sequence $(S_i)_{i\in\{0,\dots,k\}}$ corresponds to a sequence of permutations, say $(\sigma_i)_{i\in\{0,\dots,k\}}$, and the associated sequence of flips corresponds to a sequence of transposition, say $(\tau_i)_{i\in\{1,\dots,k\}}$, such that for all $i \in \{1, \dots, k\}$ $\sigma_{i-1}\tau_i = \sigma_i$.

Two transpositions do not commute unless they are identical or they operate on disjoint subsets of $\{1, \dots, n\}$. Yet, we may *swap* two transpositions in a wicker sense, preserving at least one of them. Specifically, for all distinct four integers a, b, c, d in $\{1, \dots, n\}$, we have the following identities from elementary algebra.

$$(ab)(ab) = \mathrm{Id} \tag{9.1}$$
$$(ab)(cd) = (cd)(ab) \tag{9.2}$$
$$(ab)(bc) = (ca)(ab) = (bc)(ca) \tag{9.3}$$

The idea is to transform any flip sequence (seen as a sequence of transpositions) into a flip sequence of length at most $\binom{n}{2}$ where the *track* of each pair of segments (i.e., the sequence of pairs of segments, each *turning into* the next one) is only flipped once. This may be achieved by iteratively choosing a flip f in the sequence which is the second flip of its track and swapping the transposition τ corresponding to f with the others transpositions (using (9.2) or (9.3)) to place τ next to its first occurrence in the track and cancel these to copies of τ (using (9.1)). The swapping of τ is guided by the track, i.e., choosing which of the two identities in (9.3) to use for a swap is done in order to "preserve" the states of the pairs of segments involved in the swap. It is still work in progress whether the final sequence of transposition we obtain is a flip sequence, but our experiments swapping the longest sequence known (starting with a butterfly, see Theorem 4.2.1) work well.

9.3 Untangling with Insertion Choice

It may be surprising that a quadratic upper bound is known on $\mathbf{d}^{\mathtt{I}}_{\mathtt{Multigraph}}$ (Theorem 3.1.4 [19] and Theorem 3.1.5 [17]) but not on $\mathbf{d}^{\mathtt{R}}_{\mathtt{Multigraph}}$ (where the best upper bound known is cubic in general, thanks to Theorem 3.1.3 [19, 82]). Insertion choice is indeed very different from removal choice; any simple reduction between the two seems unlikely. Next, we discuss about lower bounds, then about the recursive approach.

Lower bounds. We have not coded any experiments for insertion choice. In the following, we discuss about how to prove lower bounds nonetheless, with many aspect similar to lower bounds with removal choice.

Similarly to lower bounds with removal choice, proving a lower bound amounts to devising a strategy in place of the adversary, this time using removal choice to make the untangle sequence as long as possible. Again, the choice of the initial multiset of segments is important. Again, to fully benefit from this choice, we have to use removal choice to decrease as little as possible a potential for which we know a high lower bound on its initial value and an low upper bound on its final value. Again, the only three potentials known to have a low upper bound on crossing-free multisets of segments are the number of crossings χ and the sum, or the product, of the crossing depths $\delta_\times(s)$ of the segments s in the multiset. Recall that there are flip situations where both insertion choices lead to a big decrease of any of these three potentials. Avoiding these situations or proving that they are not frequent would lead to a lower bound on $\mathbf{d}^{\mathtt{I}}_{\mathtt{Convex\ Multigraph}}$.

As usual, the insertion strategies are not known to be optimal. Even in the `Convex` version, the algorithm from Theorem 6.1.1 providing an $O(n \log n)$ upper bound on $\mathbf{d}^{\mathtt{I}}_{\mathtt{Convex\ Multigraph}}$ is not optimal. It is uncertain whether this $O(n \log n)$ upper bound or the linear lower bound on $\mathbf{d}^{\mathtt{I}}_{\mathtt{Convex\ Multigraph}}$ can be improved. Regardless, if we set our mind to improve the lower bound, it may be a good approach to choose a multiset having each of its segments crossing all its others (with endpoints in convex position) as the starting multiset, and to use a similar removal strategy as in the proof of Theorem 5.2.1 [31]. Yet, it remains to prove that the crossing depth potentials do not decrease too much or too often.

The author's current belief is that $\mathbf{d}^{\mathtt{I}}_{\mathtt{Multigraph}}$ is $\Theta(n \log n)$ (even in the non-`Convex` version). If, regardless of this intuition, we want to prove a lower bound on $\mathbf{d}^{\mathtt{I}}_{\mathtt{Multigraph}}$ which is $\omega(n^{\frac{3}{2}})$, a consequence of Theorem 3.1.5 [17] is that, we need a point set with a crossing type that admits no realisation with an $O(\sqrt{n})$ spread. In other words, we need to find a crossing type that is impossible to draw as a dense point set.

A recursive approach. Again, a recursive insertion strategy which untangles a new segment added to a crossing-free multiset of segments in sub-linear time would lead to a sub-quadratic bound for $\mathbf{d}^{\mathtt{I}}_{\mathtt{Multigraph}}$.

9.4 Untangling with Both Choices

It is surprisingly not obvious how to take advantage of both choices to untangle segments when there is such a gap between the quadratic upper bound in general (Theorem 3.1.4 [19]) and the linear lower bound (Theorem 3.2.12). Next, we discuss about lower bounds, about the recursive approach, and about the intriguing only theorem specific to the `Matching` version.

A tight linear lower bound. As with removal choice, all our small scale experiments have so far shown that the linear lower bound seem to be asymptotically tight.

Regardless, proving a lower bound amounts to choosing an initial multiset of segments such that all its untangle sequences are long. This may be achieve similarly to the proofs of Theorem 3.2.12 and Theorem 5.1.1 using a flip invariant of some class of multiset of segments in which we have chosen the initial multiset.

A recursive approach. Once more, a recursive strategy for both choices which untangles a new segment added to a crossing-free multiset of segments in sub-linear time would lead to a sub-quadratic bound for $\mathbf{d}^{\mathtt{RI}}_{\mathtt{Multigraph}}$.

The only theorem specific to the `Matching` version. It is intriguing that Lemma 7.2.1 [31], and thereby Theorem 7.2.3 [31] are so far the only results specific to the `Matching` version which do not admit a straightforward generalization to the `Multigraph` version.

9.5 Intractability

Designing strategies to untangle segments with optimal worst case flip complexity is one thing, computing the shortest untangle sequence of each instance is another. This raises the problem of the complexity class of the latter. Theorem 8.0.1 shows that this problem is $\mathcal{NP}$-hard in the `Bipartite Matching` version. Although a direct corollary of Theorem 8.0.1 is that the shortest untangle sequence (or any constant factor approximation of it) is also $\mathcal{NP}$-hard in the `Bipartite Multigraph` version, it remains unclear how to prove intractability in any other versions. We have also consider proving the intractability of computing the longest untangle sequence of each instance. In the following, first we present ideas about why the shortest untangle sequence problem in the `Convex Bipartite Matching` version could possibly be in $\mathcal{P}$, then we explain the main difficulties we have encountered to prove the intractability of the shortest untangle sequence in other versions and of the longest untangle sequence, and we conclude by a remark about a hypothetical unified intractability proof.

The shortest untangle sequence in the `Convex Bipartite Matching` version. The proof of Theorem 8.0.1 relies on gadgets which do not have endpoints in convex position. In fact, it might not be the case that the shortest untangle sequence is $\mathcal{NP}$-hard in the `Convex Bipartite Matching` version.[2] In the following, we try to give some ideas for an algorithm to compute the shortest untangle sequence in the `Convex Bipartite Matching` version.

As in any `Bipartite` version, the segments of a bipartite matching S (with end points in convex position) are oriented. Let us explicitly represent this orientation with arrows pointing from the red point to the blue point (Figure 9.1). Recall from the exhaustive version of the proof of Lemma 3.2.3 that $\chi(S)$, the total number of crossings in S, decreases by more than one in two specific situations (both corresponding to the rightmost case in Figure 5.3). Let us call these two situations the *triangle* situations. They consists in three pairwise crossing segments. One of the triangle situations occurs when the orientation of the three segments form an oriented cycle (if we cut off the portions of segments with no crossings and containing an endpoint). Let us call it the *cyclic* triangle situation (Figure 9.1(b)). The other triangle situation occurs when only two of the three segments form an oriented path. Let us call it the *shortcut* triangle situation (Figure 9.1(e)). In a cyclic triangle situation, flipping any two of the three segments leads to an extra drop of 2 in $\chi(S)$ (for instance, one of these three flips transforms Figure 9.1(b) into Figure 9.1(a)), while in a shortcut triangle situation, the only two segments leading to such a drop are the segments of the oriented path we mentioned (this flip transforms Figure 9.1(e) into Figure 9.1(d)).

[2]This was also discussed at the workshop [86].

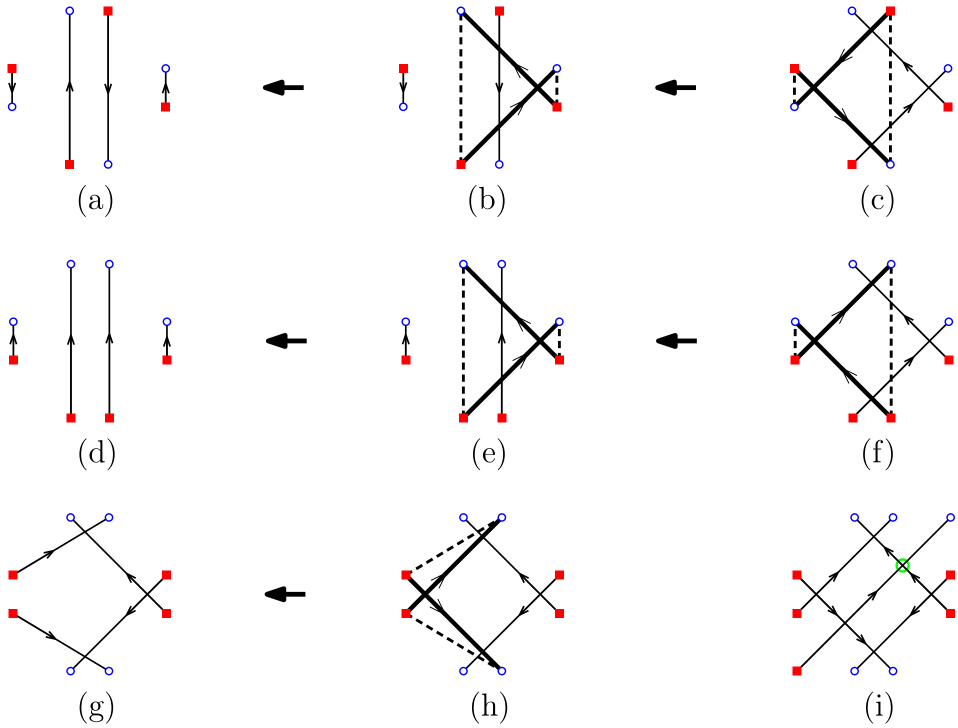

Figure 9.1: Bipartite matchings with endpoints in convex position and segments decorated with an arrow pointing towards the blue point. The sequences (c)(b)(a), (f)(e)(d), and (h)(g) are flip sequences.
(a) & (d) Crossing-free bipartite matchings.
(b) A cyclic triangle situation.
(c) & (f) Quadrilateral situations.
(e) A shortcut triangle situation.
(g) A bipartite matching without k-gon situation.
(h) A dead-end quadrilateral situation.
(i) The circled crossing is both a good choice and a bad choice: its weight is $1 - 1 = 0$.

Note that the greedy removal strategy maximizing the decrease of the potential $\chi(S)$ is not well defined if there are no triangle situation; and there exists convex bipartite matchings with no triangle situation and with untangle sequences shorter than others (Figure 9.1(c) and Figure 9.1(f) show two examples).

The difficulty is that a triangle situation may be created only by flipping two segments forming an oriented path in what we call a *quadrilateral* situation (Figure 9.1(c) and Figure 9.1(f) may be flipped into Figure 9.1(b) and Figure 9.1(e)). Again, a quadrilateral situation may only come from a *pentagon* situation using the same type of oriented path segment flipping, and so on.

Generalizing, let the *crossing* graph of a bipartite matching S be the directed graph whose vertices are crossing pairs of segments and such that there is an directed edge between two crossing pairs of segments if they share a segment (the orientation of this segment gives the orientation of the edge). Let k be a positive integer at least 3. A *k-gon* situation occurs when the crossing graph of bipartite matching S with endpoints in convex position admits a k-cycle such that non-adjacent vertices in the

k-cycle are also not adjacent in the crossing graph. If $k \geq 4$ and if the k-cycle of a k-gon situation admits a directed path of length 2, then flipping the two corresponding segments transforms the k-gon situation into a $k-1$-gon situation. We say that a k-gon situation is a *dead-end* if no flip sequence transforms it into a triangle situation (k is then even). Figure 9.1(h) shows a dead-end quadrilateral situation; an example of a flip leads to Figure 9.1(g).[3]

Obviously, our algorithm should perform in priority the flips removing pairs of segments forming an oriented path in some polygon situation which is not a dead-end; let us call this type of removal choice a *good* choice. If there is no such pair, or if there is no polygon, then it is easy to prove that any untangle sequence uses as many flips as there are crossings in S. The question is how to choose the removed pair when there is more than one good choice, or when a choice is simultaneously good in a polygon situation and not good in another polygon situation. Moreover, there are also *bad* choices, i.e., choices breaking a k-gon situation which is not a dead-end; and a choice may be simultaneously good and bad (Figure 9.1(i)). This hints that we should weight each choice with the number of polygon situations making it a good choice minus the number of polygon situations making it a bad choice (in Figure 9.1(i), the circle crossing has a weight of 0 as it is involved as a good choice in exactly 1 quadrilateral situation and as a bad choice in exactly 1 quadrilateral situation; in fact 0 appears to be the maximum weight in this matching). Is remains to prove that the strategy prioritizing the removal choice with the highest weight indeed computes one of the shortest untangle sequences, and that it runs in polynomial time using some dynamic data structure.

The shortest untangle sequence in `Redonaline Matching` version. The proof of Theorem 8.0.1 relies on gadgets which do not have colinear red points. It is thus unclear whether reducing the rectilinear planar monotone 3-SAT problem is a good approach.

The shortest untangle sequence in non-bipartite versions. The proof of Theorem 8.0.1 takes advantage of the fact that the pair of segments inserted at a flip depends only on the color of the four endpoints involved and not on all the multiset of segments. This approach is not possible in the `Cycle` version, in the `Tree` version, or in the non-bipartite `Multigraph` version.

The longest untangle sequence. We have tried to reduce the Hamiltonian path problem in grid graphs (this problem is $\mathcal{NP}$-hard [58]) to the longest untangle sequence problem (in the `Bipartite` version). We did not succeed in building a gadget which transmits the untangling chained reaction to one of the outgoing edges of a vertex. The author's current intuition is that it is geometrically impossible to build such a gadget.

A unified intractability proof. If the shortest (respectively the longest) untangle sequence of a given multiset of segments is $\mathcal{NP}$-hard to compute in all the non-`Convex`

[3] A fence only has quadrilateral situations and all of them are dead-ends. Each segment of a fence (except the leftmost and rightmost ones) participates in two quadrilateral situations. The author currently believes that this property optimizes the number of crossings while having only dead-end situations, and that is is achieved only by a fence, hinting that the lower bound provided by Theorem 5.1.1 may be tight.

versions (as we currently believe), then it is likely that there exists a proof by reduction easily adaptable to all these versions.

9.6 Conclusion

In this dissertation, we have studied the problem of untangling segments in the plane, a *reconfiguration problem* with notable applications in *TSP approximation algorithms*. We have clarified and formalized a common framework based on the notion of *choice* to study the different versions of flips combining *point set properties*, *degree properties*, and *insertion properties*. We have gathered the fourteen results about untangling segments from the literature, generalizing many of them. Throughout eighteen theorems, we have proved several *reductions* between different versions, two *counter-intuitive* lower bounds, fourteen upper bounds, and one intractability result. We have used well-known techniques such as *potentials*, balancing arguments, preprocessing, induction, and divide and conquer, as well as new techniques such as *state tracking*. We have initiated a systematic study of *point sets near convex position*.

Our problematic on untangling segments in the plane is easy to state yet hard to solve; it is deeply related to *crossing types* and *order types*, an active field of research with many essential open problems. The special version of the problem where the endpoints form a *convex* polygon is well understood and nearly solved. However, for endpoints in general position, there remain striking gaps between quadratic lower bounds and cubic upper bounds on $\mathbf{d}^{\emptyset}$, between linear lower bounds and cubic upper bonds on $\mathbf{d}^{\mathtt{R}}$, and between linear lower bounds and quadratic upper bonds on $\mathbf{d}^{\mathtt{I}}$ and $\mathbf{d}^{\mathtt{RI}}$. While we showed that the convex version does not always maximize the length of a flip sequence, the conjecture that the quadratic upper bounds on $\mathbf{d}^{\emptyset}_{\mathtt{Convex}}$, the upper bounds on $\mathbf{d}^{\mathtt{R}}_{\mathtt{Convex}}$ and $\mathbf{d}^{\mathtt{I}}_{\mathtt{Convex}}$ (both $O(n \log n)$ or linear depending on the version), and the linear upper bounds on $\mathbf{d}^{\mathtt{RI}}_{\mathtt{Convex}}$ *asymptotically* hold in the general version is still *open.*[4]

[4]The author thinks that knowledge, and maybe even truth, is more a collection of open questions than a series of answers.

Bibliography

[1] Pankaj K. Agarwal and R. Sharathkumar. Approximation algorithms for bipartite matching with metric and geometric costs. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 555–564, 2014. doi:10.1145/2591796.2591844.

[2] Oswin Aichholzer, Franz Aurenhammer, and Hannes Krasser. Enumerating order types for small point sets with applications. *Order*, 19:265–281, 2002. doi:10.1023/A:1021231927255.

[3] 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.

[4] 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.

[5] Oswin Aichholzer, Vincent Kusters, Wolfgang Mulzer, Alexander Pilz, and Manuel Wettstein. An optimal algorithm for reconstructing point set order types from radial orderings. *International Journal of Computational Geometry and Applications*, 2017. URL: https://arxiv.org/abs/1507.08080, doi:10.1142/s0218195917600044.

[6] 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.

[7] Jin Akiyama and Noga Alon. Disjoint simplices and geometric hypergraphs. In *Third international conference on Combinatorial mathematics*, pages 1–3, 1989.

[8] 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.

[9] 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.

[10] 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.

[11] Ivan J. Balaban. An optimal algorithm for finding segments intersections. *SCG '95*, 1995. doi:10.1145/220279.220302.

[12] C. Bradford Barber, David P. Dobkin, and Hannu Huhdanpaa. The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software (TOMS)*, 22(4):469–483, 1996.

[13] Jon Louis Bentley and Thomas A. Ottmann. Algorithms for reporting and counting geometric intersections. *IEEE Transactions on computers*, 9:643–647, 1979.

[14] 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.

[15] 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.

[16] Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. *Computational Geometry*. Springer, 2008.

[17] Ahmad Biniaz, Anil Maheshwari, and Michiel Smid. Flip distance to some plane configurations. *Computational Geometry*, 81:12–21, 2019. URL: https://arxiv.org/abs/1905.00791.

[18] Marthe Bonamy, Nicolas Bousquet, Marc Heinrich, Takehiro Ito, Yusuke Kobayashi, Arnaud Mary, Moritz Mühlenthaler, and Kunihiro Wasa. The perfect matching reconfiguration problem. In *44th International Symposium on Mathematical Foundations of Computer Science*, volume 138 of *LIPIcs*, pages 80:1–80:14, 2019. doi:10.4230/LIPIcs.MFCS.2019.80.

[19] Édouard Bonnet and Tillmann Miltzow. Flip distance to a non-crossing perfect matching. *arXiv*, 1601.05989, 2016. URL: http://arxiv.org/abs/1601.05989.

[20] Prosenjit Bose and Ferran Hurtado. Flips in planar graphs. *Computational Geometry*, 42(1):60–80, 2009. doi:https://doi.org/10.1016/j.comgeo.2008.04.001.

[21] 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.

[22] Maike Buchin and Bernhard Kilgus. Fréchet distance between two point sets. *Computational Geometry*, 102:101842, 2022. doi:10.1016/j.comgeo.2021.101842.

[23] 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.

[24] Bernard Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete & Computational Geometry*, 10:377–409, 1993.

[25] Bernard Chazelle and Herbert Edelsbrunner. An optimal algorithm for intersecting line segments in the plane. *[Proceedings 1988] 29th Annual Symposium on Foundations of Computer Science*, pages 590–600, 1988. doi:10.1145/147508.147511.

[26] 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.

[27] Kenneth L. Clarkson, Richard Cole, and Robert E. Tarjan. Randomized parallel algorithms for trapezoidal diagrams. In *Proceedings of the seventh annual symposium on Computational geometry*, pages 152–161, 1991.

[28] Kenneth L. Clarkson and Peter W. Shor. Applications of random sampling in computational geometry, ii. *Discrete & Computational Geometry*, 4:387–421, 1988. doi:10.1145/73393.73394.

[29] William J. Cook. *In pursuit of the traveling salesman*. Princeton University Press, 2011.

[30] 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 to appear of *Lecture Notes in Computer Science*, 2023. URL: https://arxiv.org/abs/2210.12036, doi:10.1007/978-3-031-27051-2_10.

[31] Guilherme D. da Fonseca, Yan Gerard, and Bastien Rivier. Short flip sequences to untangle segments in the plane, 2023. arXiv:2307.00853.

[32] Arun Kumar Das, Sandip Das, Guilherme D. da Fonseca, Yan Gerard, and Bastien Rivier. Complexity results on untangling red-blue matchings. In *15th Latin American Theoretical Informatics Symposium (LATIN 2022)*, volume 13568 of *Lecture Notes in Computer Science*, pages 730–745, 2022. URL: https://arxiv.org/abs/2202.11857, doi:10.1007/978-3-031-20624-5_44.

[33] 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.

[34] Donald Davendra. *Traveling salesman problem: Theory and applications*. BoD–Books on Demand, 2010.

[35] Mark de Berg, Hans L. Bodlaender, Sándor Kisfaludi-Bak, and Sudeshna Kolay. An ETH-tight exact algorithm for Euclidean TSP. *IEEE Annual Symposium on Foundations of Computer Science*, 2018. URL: https://arxiv.org/abs/1807.06933, doi:10.1109/FOCS.2018.00050.

[36] Mark de Berg, Hans L. Bodlaender, Sándor Kisfaludi-Bak, Dániel Marx, Han Hoogeveen, and Tom C. van der Zanden. A framework for exponential-time-hypothesis–tight algorithms and lower bounds in geometric intersection graphs. *SIAM Journal on Computing*, 2020. doi:10.1137/20m1320870.

[37] Mark De Berg and Amirali Khosravi. Optimal binary space partitions for segments in the plane. *International Journal of Computational Geometry & Applications*, 22(03):187–205, 2012. doi:10.1142/S0218195912500045.

[38] Rex A. Dwyer. A faster divide-and-conquer algorithm for constructing delaunay triangulations. *Algorithmica*, 2:137–151, 1987. doi:10.1007/BF01840356.

[39] Herbert Edelsbrunner, Pavel Valtr, and Emo Welzl. Cutting dense point sets in half. *null*, 1994. doi:10.1145/177424.177648.

[40] Jack Edmonds. Maximum matching and a polyhedron with 0, 1-vertices. *Journal of research of the National Bureau of Standards B*, 69(125-130):55–56, 1965.

[41] Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, 17:449–467, 1965.

[42] 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.

[43] David Eppstein, Michael T. Goodrich, and Darren Strash. Linear-time algorithms for geometric graphs with sublinearly many edge crossings. *SIAM J. Comput.*, 39:3814–3829, 2008. doi:10.1137/090759112.

[44] Paul Erdös, 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.

[45] 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.

[46] Wendell Hinkle Furry and Daniel J. Kleitman. Maximal rectilinear crossing of cycles. *Studies in Applied Mathematics*, 56(2):159–167, 1977. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/sapm1977562159, doi:https://doi.org/10.1002/sapm1977562159.

[47] Harold N. Gabow. Data structures for weighted matching and nearest common ancestors with linking. In *Proceedings of the first annual ACM-SIAM symposium on Discrete algorithms*, pages 434–443, 1990.

[48] Jacob E. Goodman, Richard Pollack, and Bernd Sturmfels. The intrinsic spread of a configuration in r d. *Journal of the American Mathematical Society*, 1990. doi:10.2307/1990931.

[49] Branko Grünbaum. Polygons in arrangements generated by n points. *Mathematics Magazine*, 46(3):113–119, 1973. doi:10.1080/0025570X.1973.11976293.

[50] Leonidas J. Guibas and Jorge Stolfi. Primitives for the manipulation of general subdivisions and the computation of voronoi. *ACM Trans. Graph.*, 4:74–123, 1985. doi:10.1145/282918.282923.

[51] Gregory Gutin and Abraham P. Punnen. *The traveling salesman problem and its variations*, volume 12. Springer Science & Business Media, 2006.

[52] 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.

[53] 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.

[54] Carmen Hernando, Michael E. Houle, and Ferran Hurtado. On local transformation of polygons with visibility properties. *Theoretical Computer Science*, 289(2):919–937, 2002.

[55] John Hershberger and Subhash Suri. Applications of a semi-dynamic convex hull algorithm. *BIT Numerical Mathematics*, 32(2):249–267, 1992.

[56] Stefan Hougardy, Fabian Zaiser, and Xianghui Zhong. The approximation ratio of the 2-opt heuristic for the metric traveling salesman problem. *Operations Research Letters*, 48(4):401–404, 2020. URL: https://arxiv.org/pdf/1909.12025.pdf, doi:https://doi.org/10.1016/j.orl.2020.05.007.

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

[58] Alon Itai, Christos H. Papadimitriou, and Jayme Luiz Szwarcfiter. Hamilton paths in grid graphs. *SIAM Journal on Computing*, 11(4):676–686, 1982.

[59] Takehiro Ito, Erik D. Demaine, Nicholas J.A. Harvey, Christos H. Papadimitriou, Martha Sideri, Ryuhei Uehara, and Yushi Uno. On the complexity of reconfiguration problems. *Theoretical Computer Science*, 412(12):1054–1065, 2011. doi:https://doi.org/10.1016/j.tcs.2010.12.005.

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

[61] David S. Johnson and Lyle A. McGeoch. The traveling salesman problem: A case study in local optimization. *Local search in combinatorial optimization*, 1(1):215–310, 1997. URL: https://techfak.uni-bielefeld.de/ags/wbski/lehre/digiSA/WS0304/IntAlg/TSPchapter.pdf.

[62] Anna R. Karlin, Nathan Klein, and Shayan Oveis Gharan. A (slightly) improved approximation algorithm for metric TSP. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 32–45, 2021. URL: https://arxiv.org/abs/2007.01409.

[63] Joseph B. Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical society*, 7(1):48–50, 1956.

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

[65] Geo Leach. Improving worst-case optimal delaunay triangulation algorithms. In *4th Canadian Conference on Computational Geometry*, volume 2, page 15. Citeseer, 1992.

[66] Chi-Yuan Lo, Jiří Matoušek, and William Steiger. Algorithms for ham-sandwich cuts. *Discrete & Computational Geometry*, 11:433–452, 1994.

[67] Chi-Yuan Lo and W. L. Steiger. An optimal time algorithm for ham-sandwich cuts in the plane. In *Proc. Canad. Conf. Comput. Geom.(CCCG 90)*, pages 5–9, 1990.

[68] 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.

[69] 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.

[70] Ketan Mulmuley. A fast planar partition algorithm. i. *[Proceedings 1988] 29th Annual Symposium on Foundations of Computer Science*, pages 580–589, 1988. doi:10.1109/SFCS.1988.21974.

[71] Naomi Nishimura. Introduction to reconfiguration. *Algorithms*, 11(4), 2018. doi:10.3390/a11040052.

[72] 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.

[73] Marcos Oropeza and Csaba D. Tóth. Reconstruction of the crossing type of a point set from the compatible exchange graph of noncrossing spanning trees. *Information Processing Letters*, 2021. doi:10.1016/j.ipl.2021.106116.

[74] Christos H. Papadimitriou. The Euclidean Traveling Salesman Problem is NP-Complete. *Theoretical Computer Science*, 4(3):237–244, 1977. doi:10.1016/0304-3975(77)90012-3.

[75] 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.

[76] 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.

[77] Michael Ian Shamos and Dan Hoey. Geometric intersection problems. *17th Annual Symposium on Foundations of Computer Science (sfcs 1976)*, 1976. doi:10.1109/sfcs.1976.16.

[78] Vera Traub, Jens Vygen, and Rico Zenklusen. Reducing path TSP to TSP. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 14–27, 2020. URL: https://arxiv.org/abs/1907.10376.

[79] Pavel Valtr. *Planar point sets with bounded ratios of distances*. PhD thesis, Freie Universität Berlin, 1994. URL: https://cir.nii.ac.jp/crid/1571135649292790272.

[80] René van Bevern and Viktoriia A. Slugina. A historical note on the 3/2-approximation algorithm for the metric traveling salesman problem. *Historia Mathematica*, 53:118–127, nov 2020. doi:10.1016/j.hm.2020.04.003.

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

[82] 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.

[83] 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.

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

[85] 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.

[86] TOPPING workshop (theory and applications of point configurations and graphs), July 2022. Barcelona. URL: https://dccg.upc.edu/topping/topping-workshop/.