Linkages in primal-dual graphs

Σπύρος Μανιάτης
Α.Μ. 201208

Τριμελής Επιτροπή:
Δημήτριος Μ. Θηλυκός, Τμ. Μαθηματικών, Ε.Κ.Π.Α.
Σταύρος Κολλιόπουλος, Τμ. Πληροφορικής και Τηλ/νιών, Ε.Κ.Π.Α.
Μιχάλης Δρακόπουλος, Τμ. Μαθηματικών, Ε.Κ.Π.Α.

Επιβλέπων:
Δημήτριος Μ. Θηλυκός, Καθηγητής του Τμήματος Μαθηματικών του Πανεπιστημίου Αθηνών

ΠΕΡΙΛΗΨΗ

Ένα από τα επιτεύγματα με τη μεγαλύτερη επιρροή στη Θεωρία Γραφημάτων υπήρξε χωρίς αμφιβολία η σειρά εργασιών "*Ελλάσσονα Γραφήματα*" των Neil Robertson και Paul D. Seymour, στην οποία, έπειτα από 23 εργασίες από το 1983 έως το 2011, κατάφεραν να αποδείξουν την εικασία του Wagner. Η εικασία αυτή λέει ότι η κλάση των μη κατευθυνόμενων γραφημάτων, μερικώς διατεταγμένων με τη σχέση ελλάσσονος γραφήματος, αποτελεί well-quasi-διάταξη ή ισοδύναμα, για κάθε κλάση γραφημάτων που είναι κλειστή ως προς ελλάσσονα υπάρχει ένα σύνολο από απαγορευμένα γραφήματα ως ελλάσσονα. Μπορεί να υποστηριχθεί ότι, δεν είναι τόσο το ίδιο το τελικό αποτέλεσμα, όσο ολόκληρη η θεωρία που αναπτύχθηκε στην πορεία που είχε, και συνεχίζει να έχει, τεράστιο αντίκτυπο τόσο στη συνδυαστική όσο και στην αλγοριθμική Θεωρία Γραφημάτων. Μία από τις κυριότερες συνεισφορές τους, η οποία κατέχει και κεντρικό ρόλο στη δουλειά τους, είναι η κατασκευή ενός αλγορίθμου που λύνει το πρόβλημα των ΔΙΑΚΕΚΡΙΜΕΝΩΝ ΜΟΝΟΠΑΤΙΩΝ σε χρόνο $f(k) \cdot n^3$, όπου k είναι το πλήθος των διακεκριμένων μονοπατιών που μας ζητείται να βρούμε. Το βασικό συστατικό της απόδειξής τους είναι η ονομαστή *τεχνική της άσχετης κορυφής* (για την οποία οι πλήρεις αποδείξεις δόθηκαν σε επόμενο μέρος της σειράς) που χρησιμοποιήθηκε ευρέως στην πορεία.

Όσο σπουδαίο κι αν αποδείχθηκε πως είναι το παραπάνω αποτέλεσμα, η συνάρτηση f που εξαρτάται από το k και εμφανίζεται στη χρονική πολυπλοκότητα του αλγορίθμου, είναι ασύλληπτα μεγάλη ακόμη και για πολύ μικρές τιμές του k. Για τον λόγο αυτόν, πολλοί ερευνητές θέλησαν να βελτιώσουν αυτήν την παραμετρική εξάρτηση από το k, είτε προσπαθώντας να απλοποιήσουν τις περίπλοκες αποδείξεις των δομικών θεωρημάτων για τη γενική περίπτωση, είτε επικεντρώνοντας την προσοχή τους σε συγκεκριμένες κλάσεις γραφημάτων των οποίων τα δομικά χαρακτηριστικά θα οδηγούσαν, ίσως, σε απλούστερες αποδείξεις και καλύτερη παραμετρική εξάρτηση. Ένα μεγάλο βήμα όσον αφορά την πρώτη κατεύθυνση (παρόλο που το φράγμα για το $f(k)$ είναι $2^{2^{2^{2^{\Omega(k)}}}}$, το οποίο είναι, σαφώς, ακόμη τεράστιο) έγινε από τους Ken-ichi Kawarabayashi και Paul Wollan στο [20]. Ένα αποφασιστικό βήμα προς τη δεύτερη κατεύθυνση, για την κλάση των *επίπεδων γραφημάτων*, έγινε από τους Isolde Adler, Stavros

G. Kolliopoulos, Philipp Klaus Krause, Daniel Lokshtanov, Saket Saurabh, και Dimitrios M. Thilikos στο [2], όπου αποδεικνύουν ένα φράγμα για το $f(k)$ που είναι απλά εκθετικό ως προς k.

Βασιζόμενοι σε αυτήν την τελευταία εργασία, μελετάμε μια επέκταση του προβλήματος των ΔΙΑΚΕΚΡΙΜΕΝΩΝ ΜΟΝΟΠΑΤΙΩΝ στην κλάση των *πρωτεύοντων-δυϊκών γραφημάτων* και χρησιμοποιώντας την ιδέα άσχετης κορυφής, αποδεικνύουμε ένα δομικό θεώρημα το οποίο λέει ότι αν το δεντροπλάτος του πρωτεύοντος-δυϊκού γραφήματος μας είναι αρκούντως μεγάλο, τότε υπάρχει (και μπορεί να εντοπιστεί αλγοριθμικά) ένα τμήμα του το οποίο είναι άσχετο και του οποίου η αφαίρεση από το γράφημα οδηγεί σε ένα απλούστερο και ισοδύναμο στιγμιότυπο του προβλήματος. Επίσης, εξηγούμε πως ένας αλγόριθμος για το πρόβλημα των ΔΙΑΚΕΚΡΙΜΕΝΩΝ ΜΟΝΟΠΑΤΙΩΝ για την κλάση των πρωτεύοντων-δυϊκών γραφημάτων μπορεί να χρησιμοποιηθεί για την κατασκευή αλγορίθμων για προβλήματα σε ενεπίπεδα γραφήματα, όπου είναι απαραίτητο να λαμβάνεται υπόψην η τοπολογία της επίπεδης εμβάπτισης που δίνεται ως είσοδος.

ABSTRACT

One of the most influential bodies of work in Graph Theory has, undoubtedly, been the *Graph Minor series* of Neil Robertson and Paul D. Seymour, where, after 23 papers during the years 1983-2011, they managed to prove *Wagner's conjecture*. This conjecture states that undirected graphs, partially ordered by the graph minor relationship, form a well-quasi-ordering, or, equivalently, every family of graphs that is closed under minors can be defined by a finite set of forbidden minors. One can argue that it is not just the final result itself, but whole theory built during the procedure which had, and continues to have, a huge impact in both combinatorial and algorithmic Graph Theory. One of their main contributions, which also has a central role in their work, is constructing an algorithm that solves the Disjoint Paths problem in $f(k) \cdot n^3$ steps, where k is the number of disjoint paths that we are asked to find. The key ingredient of their proof is the so called *irrelevant-vertex technique* (for which full proofs only appeared in latter parts of the series), which has been used extensively thereafter.

As great as this result was proved to be, the function f of k that appears in the running time is immense even for very small values of k. Therefore, many researchers tried to improve this parametric dependance on k, either by trying to simplify the complicated proofs of the structural theorems for the general case, or by restricting their attention to specific graph classes whose structural characteristics would hopefully lead to simpler proofs and better parametric dependance. A big step towards the first direction (although the bound of $f(k)$ is $2^{2^{2^{2^{\Omega(k)}}}}$ which is of course still huge) was made by Ken-ichi Kawarabayashi and Paul Wollan in [20]. A decisive step to the second direction, for the class of *planar graphs*, was made by Isolde Adler, Stavros G. Kolliopoulos, Philipp Klaus Krause, Daniel Lokshtanov, Saket Saurabh, Dimitrios M. Thilikos in [2], where their bound for $f(k)$ is just single exponential on k.

Based on this latter work, we study an extension of the Disjoint Paths problem for the class of *pd-graphs* and, using on the idea of the irrelevant-vertex technique, we prove a structural theorem which states that if the treewidth of our pd-graph is sufficiently large, then there exists (and can be found algorith-

mically) a part of it which is irrelevant and whose removal leads to a simpler and equivalent instance. We also illustrate how an algorithm for the Disjoint Paths problem for the class of pd-graphs can be used to construct algorithms for problems on plane graphs, where the it is essential to respect the topology of the plane embedding given as an input.

ΠΡΟΛΟΓΟΣ

Αυτή η διπλωματική εργασία εκπονήθηκε στα πλαίσια της ολοκλήρωσης των σπουδών μου στο *Μεταπτυχιακό Πρόγραμμα Λογική, Αλγορίθμων και Θεωρίας Υπολογισμού (Μ.Π.Λ.Α.).*

Θα ήθελα να ευχαριστήσω ειλικρινά κάποιους ανθρώπους, αν και ελπίζω να μου δοθεί η ευκαρία να τους δείξω την ευγνωμοσύνη μου στην πράξη και όχι μόνο στα λόγια αυτού του προλόγου.

Ευχαριστώ ιδιαίτερα τον επιβλέποντα της παρούσας εργασίας, Καθηγητή κ. Δημήτριο Θηλυκό, για τις ευκαιρίες που μου έδωσε, τη στήριξη που μου πρόσφερε όταν τη χρειάστηκα και για ό,τι έχω αποκομίσει από τις συζητήσεις μας, μαθηματικού ή όχι περιεχομένου.

Ευχαριστώ τον Αναπληρωτή Καθηγητή κ. Σταύρο Κολλιόπουλο για την τιμή που μου έκανε να είναι μέλος της τριμελούς επιτροπής μου και γιατί έχει υπάρξει για μένα πρότυπο ακαδημαϊκού, και όχι μόνο, ήθους.

Ευχαριστώ τον Αναπληρωτή Καθηγητή κ. Μιχάλη Δρακόπουλο που με τιμά με τη συμμετοχή του στην τριμελή επιτροπή αυτής της εργασίας.

Ευχαριστώ από καρδιάς τους συμφοιτητές μου για όλα τα όμορφα και τα άσχημα που μοιραστήκαμε.

Πάνω απ' όλα ευχαριστώ τους δικούς μου ανθρώπους για την αγάπη τους και την πάντα αδιαπραγμάτευτη στήριξή τους.

Σπύρος Μανιάτης
Αθήνα, 3 Ιουνίου 2016

CONTENTS

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

In this first section we describe the subject of this thesis, although some of the notions used here are only defined later. The reader that is not familiar with the basics of Computational Complexity and Graph Theory, is advised to postpone the reading of this introduction until the end of Chapter 3. We first present the DISJOINT PATHS problem, along with some important results, and then we briefly discuss its connection with our subject.

1.1 The Disjoint Paths Problem

One of the most central problems in graph theory, from the algorithmic point of view, is to efficiently decide whether two nodes v, u of a graph G belong to the same connected component of G, i.e. if there exists some path of G with endpoints v and u. To state it more formally:

REACHABILITY
Input: A graph $G = (V, E)$, and two vertices $v, u \in V$.
Question: Is there a path P with endpoints v and u in G?

It is widely known that the REACHABILITY problem admits polynomial time algorithms (depth-first search, breadth first search). One can wonder though what happens if we wish to determine whether there exist paths between multiple pairs of vertices in a given graph G. Of course if we allow our paths to intersect on edges, then an answer can be obtained in polynomial time using one of the algorithms that we already now for solving REACHABILITY. But what happens if we demand our paths to be edge-disjoint or even vertex-disjoint, meaning that two or more paths cannot share an edge or vertex, respectively? In this work we will focus on the vertex-disjoint version of the problem:

DISJOINT PATHS (DP) **Input**: An undirected graph $G = (V, E)$, and a set $\mathcal{P} = \{(s_i, t_i) \mid i = 1, \ldots, k\}$ of pairs of vertices of G. **Question**: Do there exist paths $P_1, \ldots, P_k$ of G which are mutually vertex-disjoint and such that the endpoints of P_i are s_i and t_i, for every $i \in [k]$?

We will sometimes refer to the collection of pairs of vertices $\mathcal{P}$, as the *pattern* of the input and to the vertices of its pairs, as *terminals*.

DP was shown to be NP-complete by Karp in [16] and the same holds even if the input graph is restricted to be planar as proven in [41]. NP-completeness is also the case for the edge-disjoint and directed variants of DP, as indicated by [25] and [22].

What happens if we are given a really large graph, i.e. n is very big, but we are asked to find two disjoint paths connecting two given pairs of vertices? The question translates to whether there exists some efficient (polynomial?) algorithm for solving DP when k is fixed to be two and the answer was given in 1980 with the polynomial algorithms presented independently in [36], [37] and [39].

But as almost always happens natural questions keep coming: What is the best we can do if the number of pairs we want to join with a path is fixed, meaning that it is not part of the input but is given as a parameter. Or, if we want to express it in terms of Parameterized Complexity, what is the complexity of the following parameterized problem?

p-DISJOINT PATHS (p-DP) **Input**: An undirected graph $G = (V, E)$. **Parameter**: Positive integer k and a set $\mathcal{P} = \{(s_i, t_i) \mid i = 1, \ldots, k\}$ of k pairs of vertices of G. **Question**: Do there exist paths $P_1, \ldots, P_k$ of G which are mutually vertex-disjoint and such that the endpoints of P_i are s_i and t_i, for every $i \in [k]$?

The answer was given by Robertson and Seymour in the 13th part of their Graph Minor series of papers, [31], as they presented an algorithm that solves p-DP in $f(k) \cdot n^3$ steps, where f is some computable function, thus classifying p-DP in FPT (actually they considered the DP problem as the framework of parameterized complexity was not explicit at that time). This algorithm has a central role in their work, which ends up to proving the Wagner's conjecture [32], which is considered to be one of the greatest and most influential achievements of graph theory in the last decades.

The algorithm in [31] is based on the *irrelevant-vertex technique*, which developed by Robertson and Seymour and used widely ever since in many combinatorial problems (see for example [7], [8], [13], [17], [18], [19] and, [14]) and is based on the following idea when applied to a problem Π on graphs:

As long as the input graph G violates some specified structural conditions, there exists (and can be found efficiently), a *solution-irrelevant* vertex, i.e. a vertex of G whose deletion does not result to a different answer to the question imposed in problem Π. One then iteratively locates and removes such irrelevant vertices until the structural conditions are met, at which point the graph has been simplified enough and the problem can be solved using known tools.

For the case of DP problem, an irrelevant vertex is a vertex of the input graph G with the following property: Any collection of paths in G that certifies a solution can be transformed into an *equivalent* one (meaning that it links the same pairs of vertices) in $G \setminus \{v\}$, which roughly suggests that vertex v is not necessary to link the given pairs of vertices and thus can be discarded. The structural conditions used by the algorithm in [31] are the following two:

(i) G excludes a clique, of a certain size which depends on k, as a minor.

(ii) The treewidth of G is bounded by a function, say g, of k.

The most complicated part of their proof, on which the correction of their algorithm heavily relies, and which was postponed until the later papers [33], [34], was to show that when Condition (i) is met, i.e. graph G does not contain any "big" clique-minors, if $\mathbf{tw}(G) \geq g(k)$, then there exists an irrelevant vertex. The drawback of this algorithm is that the parametric dependance on k, expressed by $f(k)$, is huge due to the bounds that arise from the complicated proofs, making it almost useless for practical purposes (and remember that one of the main motivations of parameterized algorithms is to fight intractability through fine-grained analysis and ultimately to be able to solve an NP-complete (or harder) problem efficiently (in practical and terms) when some parameters are bounded). Consequently, this need to greatly improve the parametric dependance that emerges from the structural theorems in the Graph Minor series, lead the researchers to either try to simplify parts of the proofs in the series (as in [20] where a $2^{2^{2^{\Omega(k)}}}$, still pretty huge, lower bound is achieved for $f(k)$) or to restrict their attention to specific graph classes whose structural characteristics will hopefully lead to simpler proofs and better parametric dependance.

1.2 About this thesis

In this thesis we study an extension of the DP problem in the class of pd-graphs, which will be defined formally in the next section, thus working in the context of the second direction as described previously.

We are based on the work of Adler et al. [2]. Actually some of our proofs are similar to the ones given in [2], but we nevertheless present all proofs here in order to provide a complete picture. The problem that they study is the following:

PLANAR DISJOINT PATHS (PDP) **Input**: An undirected planar graph $G = (V, E)$ and a set P=$\{(s_i, t_i) \mid i = 1, \ldots, k\}$ of k pairs of vertices of G. **Question**: Do there exist paths $P_1, \ldots, P_k$ of G which are mutually vertex-disjoint and such that the endpoints of P_i are s_i and t_i, for every $i \in [k]$?

The main results of their work are the followings:

Proposition 1.2.1. *Every instance of* PDP *consisting of a planar graph G of treewidth at least $\mathcal{O}(2^k)$ and k pairs of terminals, contains a vertex v such that every solution to PDP can be replaced by an equivalent one whose paths avoid v.*

This result has a structural essence and states that if a planar graph has "big enough" treewidth then there certainly exists some vertex v of G that can be avoided (actually they prove that not only there exists some solution that avoids v, but any solution can be transformed into one that avoids v).

Proposition 1.2.2. *There exists an algorithm that, given an instance $(G, \mathcal{P})$ of* DP*, where G is an n-vertex graph and $|\mathcal{P}| = k$, either reports that $(G, \mathcal{P})$ is a NO-instance or outputs a solution of PDP for $(G, \mathcal{P})$. This algorithm runs in $O(k) \cdot n^2$ steps.*

Their second result is algorithmic and basically turns the irrelevant vertex technique into an algorithm (as described briefly previously): If the treewidth of the input graph is "big enough", the algorithm finds an irrelevant vertex (which is guaranteed to exist as the structural condition in this case is exactly the treewidth being "big enough") and removes it. This is done iteratively until the treewidth of the graph becomes sufficiently small and the problem is attacked directly using already known algorithms.

Our problem. Before proceeding to the next chapter where we develop all its needed in order to state our problem, we give a brief description. A pd-graph can be though of as a plane graph and its dual graph considered as one embedding where new vertices are introduced at each intersection of edges of the initial graph and its dual (in order for our structure to be plane). Then, the vertex set of our graph is naturally partitioned into three sets: the *primal vertices* (which correspond to the vertices of the initial graph), the *dual vertices* (which correspond to the vertices of the dual) and the *crossing vertices* (which correspond to the vertices introduced at the intersections). See Figure 1.2 for a simple example.

The reason why we are interested in this kind of embedded structures and not just planar graphs as combinatorial structures (where an embedding is not necessarily unique and therefore one cannot talk about the topology of a planar graph), is because we wish to study problems that are also related to the topology of a plane graph. For example, someone could be interested in solving problems where the embedding Γ of a planar graph, i.e. a planar graph, is given and the task is to "find something inside (or outside) a face of Γ". It is clear that such kind of questions are not well defined in the context of planar graphs. The notion of pd-graph is used to "translate" the topological properties of a plane graph into combinatorial ones in order to unlock the rich toolbox of Combinatorial Graph Theory and address questions of topological nature.

Then, the variant of the DP problem that we study differs on the following sense: Each pair of the collection $\mathcal{P}$ of the input is either a *primal pair* or a *dual pair* and and if $(s_i, t_i) \in \mathcal{P}$ is a primal (resp. dual) pair we demand that in a solution the path with endpoints s_i and t_i does not contain any dual (resp. primal) vertices. Having this problem in mind, we prove some structural results leading to the existence of some irrelevant part of the input graph, given that it satisfies some structural condition.

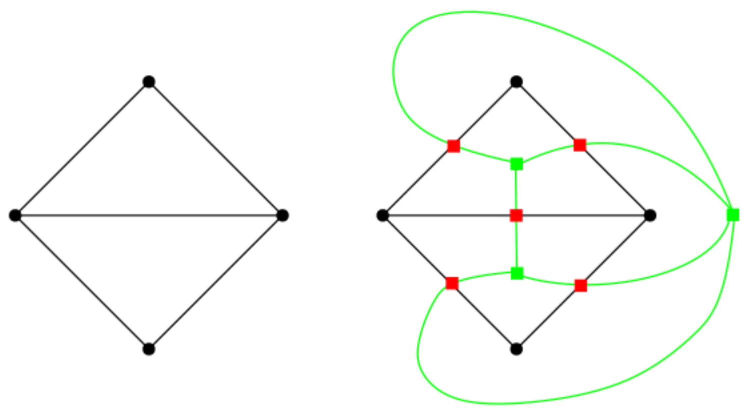

Figure 1.1: A plane graph at the left and a corresponding pd-graph at the right. The black dots correspond to the primal vertices, the green squares to the dual vertices, and the red squares to the crossing vertices.

CHAPTER 2

BASIC NOTATION

This second chapter contains some basic notation used throughout this thesis as well as some brief introduction to graphs (for a more detailed presentation see [9]) and some width parameters on graphs.

2.1 Basics

$\mathbb{N}, \mathbb{Z}, \mathbb{Q}$ and $\mathbb{R}$ denote the sets of natural numbers, integer numbers, rational numbers and real numbers, respectively. For every $n \in N$, we define $[n]$ to be the set $\{1, 2 \ldots, n\}$. Let S be a set, we denote by 2^S the set of all subsets of S and for every $k \in \{0, 1 \ldots, |S|\}$ we denote by 2^{S_k} the set of all subsets of S with exactly k elements.

2.2 Graphs

Graphs. A (*simple*) *graph* $G = (V, E)$ is a pair of sets, V and E, where $E \subseteq \{\{v, u\} \in 2^V \mid u \neq v\}$. The elements of V and E are called the *vertices* and the *edges* of the graph G, respectively. Given a graph G, we denote by $V(G)$ and $E(G)$ the vertices and edges of G, respectively.
A structure $G = (V, E)$ is called a *multigraph* if E is allowed to be a multiset and can contain $\{v, v\}$ for some $v \in V$ (such an element is called a *loop*).
A *directed graph* $G = (V, A)$ is a pair of sets, V and A, where $A =\subseteq \{\{v, u\} \in V \times V \mid u \neq v\}$. The elements of A are called *arcs* of G and given an arc $a = (u, v) \in A$, u and v are called the *head* and the *tail* of a, respectively.

Operations on graphs. Let $G = (V, E)$ be a graph and let $v \in V$ be a vertex of G. We say that $N_G(v) = \{u \in V : \{v, u\} \in E\}$ is the *open neighborhood* of v in G and $N_G(v) = \{u \in V : \{v, u\} \in E\}$ the *closed neighborhood* of v in G. We define the graph obtained from G if we *delete* vertex v, as the graph whose

vertex set is $V \setminus \{v\}$ and edge set is $E \setminus \{\{v, u\} : u \in N_G(v)\}$ and denote this graph by $G \setminus \{v\}$ (or just $G \setminus v$). Let $A \subset V$, we denote by $G \setminus A$ the graph obtained by deleting all vertices in A from G.
We define the graph obtained from G if we *delete* edge $e \in E$, as the graph whose vertex set is V and edge set is $E \setminus \{e\}$ and denote this graph by $G \setminus \{e\}$ (or just $G \setminus e$). Let $F \subset E$, we denote by $G \setminus F$ the graph obtained by deleting all edges in F from G.

Subgraphs and induced subgraphs. Let $G = (V, E)$ be a graph. We say that a graph H is a *subgraph* of the graph G if H can be obtained from G after a sequence of vertex and edge deletions (for an example see Figure 1.1).
Let $A \subseteq V$. We denote by $G[A]$ the subgraph of G whose vertex set is A and edge set is $E(G[A]) = \{e = \{u, v\} \in E : \{u, v\} \subseteq A\}$. We call $G[A]$ the subgraph of G that is *induced* by the vertices in $A \subseteq V$ (for an example see Figure 2.2).

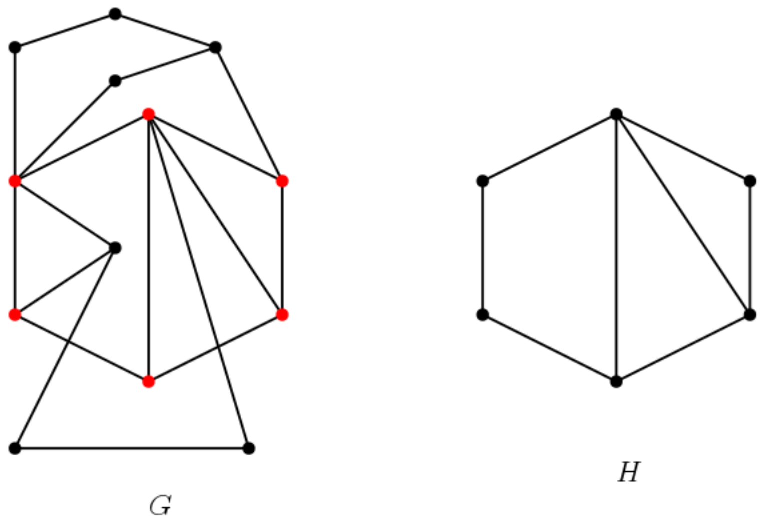

Figure 2.1: A graph G (at the left) and the subgraph H induced by the vertices that are colored red in G (at the right).

Paths and cycles. Let $G = (V, E)$ be a graph. A *path* P in G is a subgraph of G whose vertex set is some subset $\{v_1, v_2, \ldots, v_k\}$ of V and whose edge set is $\{\{v_i, v_{i+1}\} : 1 \leq i \leq k-1\}$. If $k = 1$ we say that P is a *trivial* path and when $k \geq 1$ we define the *length* of P to be $k - 1$. The vertices v_1 and v_k are called the endpoints of P.
A *cycle* C in G is a subgraph of G whose vertex set is some subset $\{v_1, v_2, \ldots, v_k\}$ of V and whose edge set is $\{\{v_i, v_{i+1}\} : 1 \leq i \leq k-1\} \cup \{v_1, v_k\}$. If $k = 1$ we say that C is a *trivial* cycle and when $k \geq 1$ we define the *length* of C to be k.
We say that the graph G is *connected* if for every two vertices $u, v \in V$, there exists a path in G with endpoints v and u. Let $k \in \mathbb{N}$, we say that the graph G

is k-*connected* if for any $A \subset V$ with $|A| \leq k-1$, the graph $G \setminus A$ is connected.

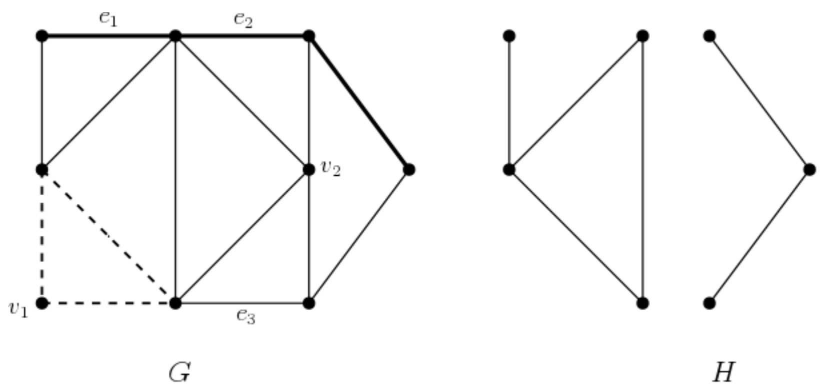

Figure 2.2: At the right: A path of length 3 in graph G (bold edges). A cycle of length 3 (or a triangle) (dashed edges). At the left: The subgraph H that results from G by deleting the vertices v_1 and v_2 and the edges e_1, e_2 and e_3.

Minors. A graph H is a *minor* of a graph G, if there exists a function $\phi : V(H) \to 2^{V(G)}$ such that

1. For every $u, v \in V(H)$ with $u \neq v$, $G[\phi(u)]$ and $G[\phi(v)]$ are two vertex-disjoint connected subgraphs of G.
2. For every edge $e = \{u, v\} \in E(H)$, $G[\phi(u) \cup \phi(v)]$ is a connected subgraph of G.

Topological minors. We say that a graph H is a *topological minor* of a graph G if there exists an injective function $\phi_0 : V(H) \to V(G)$ and a function $\phi_1 : E(H) \to \mathcal{P}(G)$ such that

- for every edge $\{x, y\} \in E(H)$, $\phi_1(\{x, y\})$ is a path between $\phi_0(x)$ and $\phi_0(y)$.
- if two paths in $\phi_1(E(H))$ have a common vertex, then this vertex should be an endpoint of both paths.

Given the pair (ϕ_0, ϕ_1), we say that H is a *topological minor of G via* (ϕ_0, ϕ_1).

Planar, plane, and outerplanar graphs. A graph G is called *planar* if it can be embedded in the plane $\mathbb{R}^2$ (or equivalently in the sphere $S^2 = \{(x, y, z) \in \mathbb{R}^3 : x + y + z = 1\}$) in such a way that there are no two edges of it whose embeddings intersect (they can meet only at their endpoints). Such an embedding is called a *planar embedding of G* and we say that such it is a *plane*

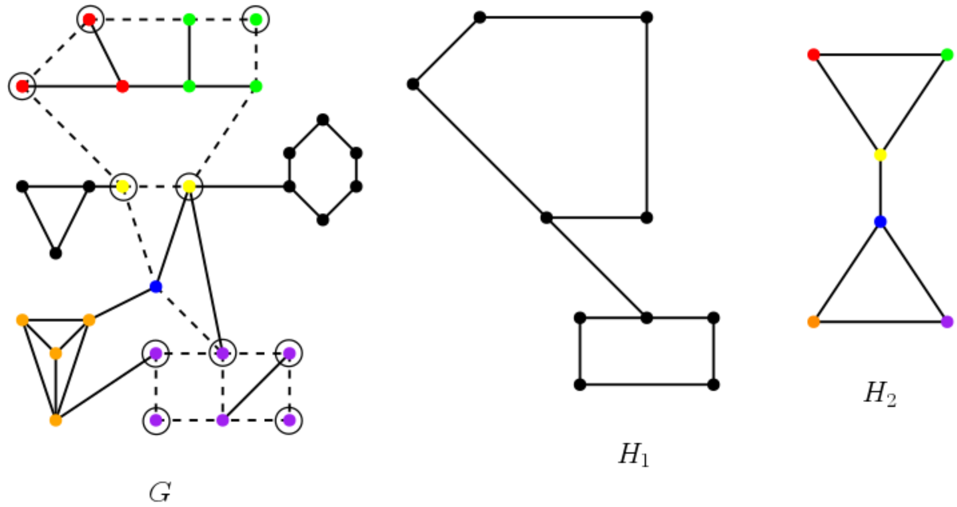

Figure 2.3: The graph H_1 is a topological minor of the graph G (certified by the circled vertices of G and the dashed edges of G) and the graph H_2 is a minor of G (consider the function $\phi : V(H) \to 2^{V(G)}$ that sends a vertex of H_2 to the subset of vertices of G of the same color and observe that each "color class" in G induces a connected subgraph.

graph (observe that a planar graph can have more than one planar embeddings that can also be different from a topological point of view). Given a plane graph G we denote its faces by $F(G)$, i.e. $F(G)$ is the set of the connected components of $\mathbb{R}^2 \setminus G$ (in the operation $\mathbb{R}^2 \setminus G$ we treat G as the set of points of $\mathbb{R}^2$ corresponding to its vertices and its edges).
The *dual*, G^*, of a plane (planar) graph G is also a plane (planar) graph and has one vertex for each face of G. There is an edge between two vertices of G^* if and only if the boundaries of their corresponding faces share an edge (observe that if a plane graph is not connected it can have, two or more, different (from a topological point of view) dual graphs). For an example of a plane graph and its corresponding dual graph see Figure 2.2.
An *outerplanar* graph is a plane graph whose vertices are all incident to the infinite face. If an edge of an outerplanar graph is incident to its infinite face then we call it *external*, otherwise we call it *internal*. The *weak dual* of an outerplanar graph G is the graph obtained from the dual of G after removing the vertex corresponding to the infinite face of the embedding. We call a face of an outerplanar graph *simplicial* if it corresponds to a leaf of the graph's weak dual. For an example see Figure 2.2.

2.3 Width parameters

Treewidth. A *tree decomposition* of a graph G is a pair (T, χ), consisting of tree T and a mapping $\chi\colon V(T) \to 2^{V(G)}$, such that for each $v \in V(G)$ there

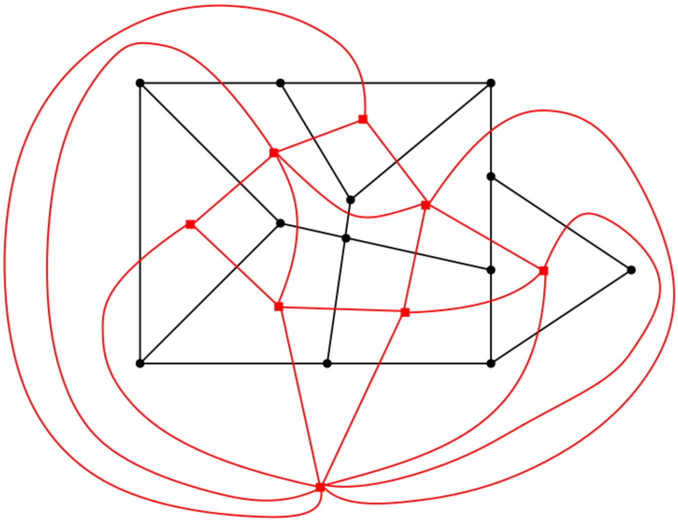

Figure 2.4: A plane graph (black) embedded in the plane along with its dual graph (red). There is one dual vertex (red square) for every face of the plane graph. Any edge of the black graph is on the boundary of exactly two of its faces which are connected by an edge in the dual (red) graph.

exists $t \in V(T)$ with $v \in \chi(t)$, for each edge $e \in E(G)$ there exists a node $t \in V(T)$ with $e \subseteq \chi(t)$, and for each $v \in V(G)$ the set $\{t \in V(T) \mid v \in \chi(t)\}$ is connected in T. (for an example see Figure 2.3). The *width* of (T, χ) is defined as **width**(T, χ) := max $\{|\chi(t)| - 1 \mid t \in V(T)\}$. and the *tree-width of G* is defined as

$$\textbf{tw}(G) := \min \{\textbf{width}(T, \chi) \mid (T, \chi) \text{ is a tree decomposition of } G\}.$$

Grids. Let $m, n \geq 1$. The *$(m \times n)$-grid* is the Cartesian product of a path of length $m - 1$ and a path of length $n - 1$. In the case of a *square grid* where $m = n$, we say that n is the *size* of the grid. Given that $n, m \geq 2$, the *corners* of an *$(m \times n)$-grid* are its vertices of degree 2. When we refer to a *$(m \times n)$-grid* we will always assume an orthogonal orientation of it that classifies its corners to the *upper left*, *upper right*, *down right*, and *down left* corner of it.

Given that Γ is an *$(m \times n)$-grid*, we say that a vertex of G is one of its *centers* if its distance from the set of its corners is the maximum possible. Observe that a square grid of even size has exactly 4 centers. We also consider an $(m \times n)$-grid embedded in the plane so that, if it has more than 2 faces then the infinite one is not a square. The *outer cycle* of an embedding of a $(m \times n)$-grid is the one that is the boundary of its infinite face. We also refer to the *horizontal* and the *vertical lines* of a *$(m \times n)$-grid* as its paths between vertices of degree smaller than 4 that are traversing it either "horizontally'' or "vertically'' respectively. We make the convention that an $(m \times n)$-grid contains m vertical lines and n horizontal lines. The *lower horizontal line* and the *higher horizontal line* of Γ are

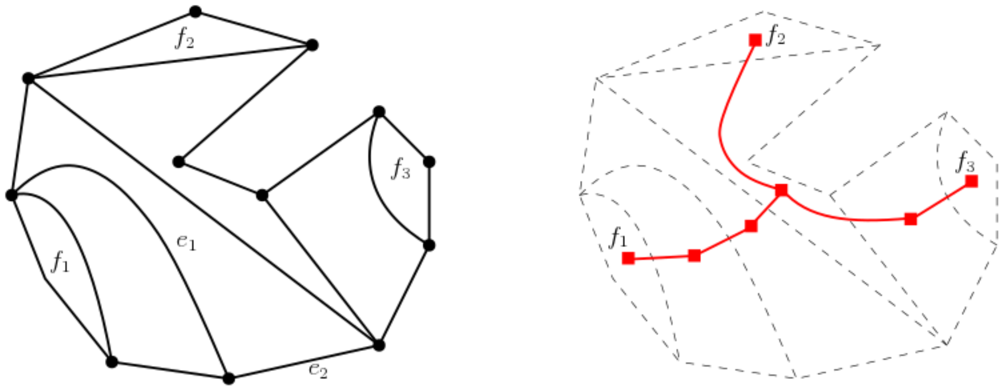

Figure 2.5: An outerplanar graph at the left and its weak dual at the right. Its simplicial faces are f_1, f_2 and f_3, e_1 is an internal edge, and e_2 is an external edge.

defined in the obvious way (for an example see Figure 2.3).

Branchwidth. A *branch decomposition* of a graph G is a pair (T, τ), where T is a tree with vertices of degree one or three and τ is a bijection from $E(G)$ to the set of leaves of T. The *order function* $\omega : E(T) \to 2^{V(G)}$ of a branch decomposition maps every edge e of T to a subset of vertices $\omega(e) \subseteq V(G)$ as follows. The set $\omega(e)$ consists of all vertices $v \in V(G)$ such that there exist edges $f_1, f_2 \in E(G)$ with $v \in f_1 \cap f_2$, and such that the leaves $\tau(f_1)$, $\tau(f_2)$ are in different components of $T - \{e\}$. The *width* of (T, τ) is equal to $\max_{e \in E(T)} |\omega(e)|$ and the *branchwidth* of G is the minimum width over all branch decompositions of G.

We will now state a proposition that follows directly by combining the next two results:

Result 1.([15]) If G is a planar graph and $\mathbf{bw}(G) \geq 3k + 1$, then G contains a $(k \times k)$-grid as a minor.

Result 2.([30]) If G is a graph, then $\mathbf{bw}(G) \leq \mathbf{tw}(G) + 1 \leq \frac{3}{2} \cdot \mathbf{bw}(G)$.

Proposition 2.3.1. *If G is a planar graph and* $\mathbf{tw}(G) \geq 4.5 \cdot k + 1$*, then* G *contains a* $(k \times k)$*-grid as a minor.*

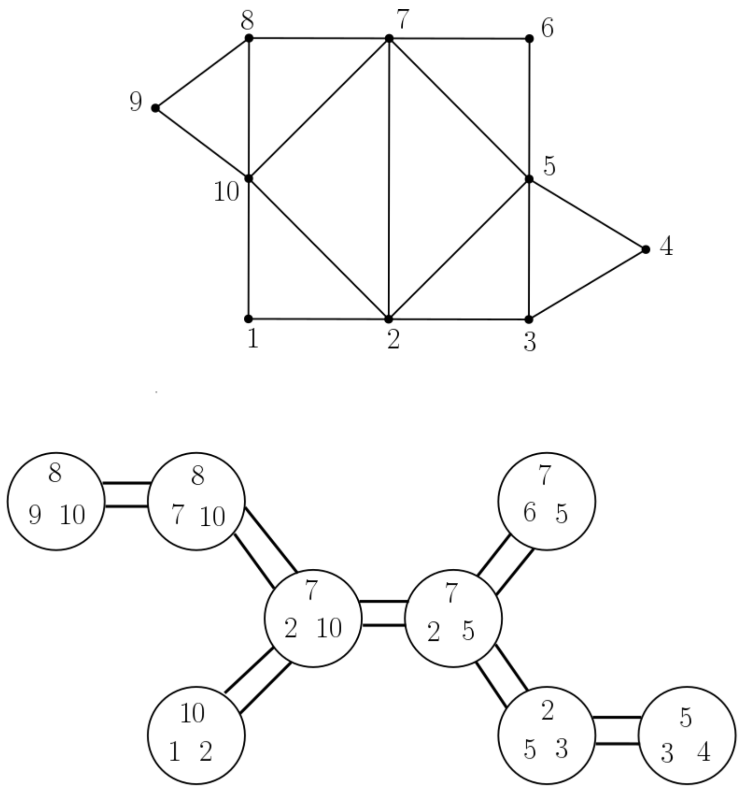

Figure 2.6: At the top there is a graph on 10 vertices and at the bottom a tree decomposition of it with width 3. It is easy to confirm that any tree decomposition of this graph has width at least 3, thus the treewidth of the depicted graph is 2.

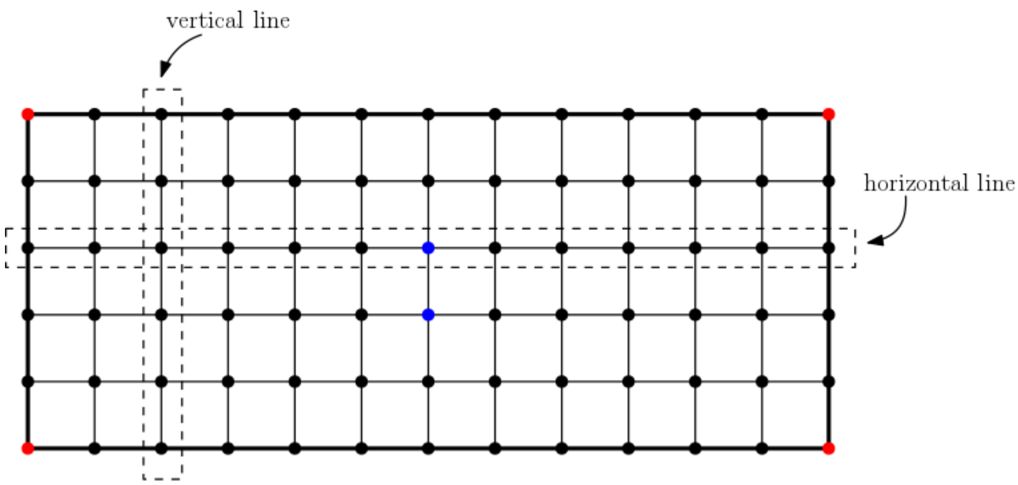

Figure 2.7: A (13×6)-grid is depicted. Its corners are the red vertices and its centers are the two blue vertices. The outer cycle is the bold rectangle that contains the corners of the grid.

CHAPTER 3

FUNDAMENTALS OF ALGORITHMS AND COMPLEXITY

3.1 Brief history of Theoretical Computer Science

In this section we attempt to present a brief overview of the history of Computational Complexity in order to define some complexity classes related to the topic of this thesis.

An important question in Mathematics, that appeared throughout their history, is whether for a given problem there exists a way for solving it that can be clearly formulated and consists of "simple" steps or computations. This is what we think of a problem that can be solved "algorithmically". Someone could wonder whether there provably exist problems that do not have the aforementioned property. Mathematicians studied these notions and developed a branch of Mathematics that is called *Computability Theory* and which gave birth to *Theoretical Computer Science* as we know it today.
But what does "algorithmically" mean? In order to properly define this notion, one needs to fix a *model of computation* which, intuitively, specifies the capabilities of the "machine" that executes our algorithms. The most widely used model, especially for theoretical purposes, is the *Turing machine*. As the topic of computational models is out of the scope of this thesis, the reader is reffered to [23] for an extensive introduction. `

The birth of computers and the idea of "machines" being able to execute long and complicated computations for us, introduced a new parameter in the Theory of Computation:

Efficiency!

The notion of efficiency was central in the analysis of computational problems and initiated the research for algorithms that achieve better performance than the existing ones. The two main measures of efficiency are *space* and *time*, and in this thesis we will focus on the latter. Given these new parameters, computational problems can be classified further in *complexity classes* based on the efficiency of algorithms that solve them. Undoubtedly, the most well known complexity classes are P and NP, where P contains the problems that can be solved efficiently (there exist deterministic algorithms that solve them in time that is bounded by a fixed polynomial on the size of the input) and NP contains the problems which require nondeterminism in order to be solved in polynomial time. It is widely believed that P $\neq$ NP, which can be roughly translated to the fact that there exist computational problems for which any algorithm solving them needs exponential time. Consequently, most computer scientists face at some point the following question when studying a computational problem:

Is there an efficient algorithm solving the problem of interest? If not, is it possible to provide some evidence that it cannot be solved efficiently?

The natural approach to address these questions is either trying to come up with a polynomial time algorithm that solves the problem (which places it in P) or proving that, assuming P $\neq$ NP, it is in NP but not in P. This can be done by *reducing* an NP-*complete* problem to the problem of interest. Roughly speaking, NP-*complete* problems are the hardest in the class NP and the reduction of such a problem to another problem suggests that the latter is at least as hard, thus characterized as intractable. For more information about the theory of NP-completeness we refer the reader to [12].

The construction of a polynomial time algorithm is usually the best outcome one can hope for (recently this becomes more and more inaccurate as we need to solve problems where the input is huge; the running time of a polynomial n^4 algorithm when the input is the web network does not seem appealing at all! In many cases even a linear algorithm can be practically useless and this means that the desired algorithm does cannot access or store all of its input. For more information on the subject we refer the reader to [35]. But what happens if we prove that our problem is NP-complete? Is this the end of the story? Fortunately, the answer is no and we briefly present the main side roads one can choose from:

- **Approximation:** A very important class of problems is the one of *optimization problems* where the task is to find the *best solution* from all feasible solutions. Unfortunately, many optimization problems have proved to be NP-complete. When the need for an exact solution is not imperative, a way to overcome this difficulty is trying to design efficient algorithms that find a solution which is guaranteed to be"close" to the optimal. Of course an analogue of intractability arises in this setting too and much work has been done in the direction of proving *innaproximability* and *lower bounds* results. Approximation algorithms have been developed rapidly in the last decades and proved to be a very fruitful area. For an extensive introduction we refer the reader to [42] and [40].

- **Use of randomness:** Another tool that can be used to cope with an NP-complete problem is *randomness* and the study of randomized algorithms was spurred by the discovery of a randomized primality test

[38]. The main idea of this approach is roughly the following: In order to "prune" some of the branches of computation, which seem to be unavoidably exponential (under worst-case analysis) when trying to solve an NP-complete problem, the randomized algorithm makes some random choices and based on them, and probably other deterministic computation, produces an answer. One has to distinguish between algorithms that use randomness in order to reduce the expected running time and always terminate in bounded time producing the right answer (called *Las Vegas algorithms*) and algorithms that terminate in polynomial time but there is a chance that they produce a wrong answer or no answer at all (called *Monte Carlo algorithms*). Having designed a randomized algorithm for a problem, it is sometimes possible to produce a deterministic algorithm solving the same problem. This procedure is known as *derandomization* and has attracted much attention recently. More information about randomized algorithms can be found in [27] and [26] and for some information about the complexity classes that arise from randomized algorithms see [24] and [3].

- **Parameterization:** When a problem is NP-complete, any exact deterministic algorithm that solves it takes (in the worst case) exponential (or at least superpolynomial) to n time, where n is the length of the input. The parameterized complexity point of view examines whether this exponential explosion on the running time unavoidably "spreads" to a large part of the input (meaning a part whose length depends on n) or there are some particular *parameters* of the problem which cause the increase on the running time. For some NP-complete problems that are of great in importance in other areas, such as biology, there were algorithms that, although being exponential in the worst case, worked efficiently in practice. Then a natural question arose:

  Are there some parameters in these particular problems which happen to be bounded and this way "soften" the intractability? Can theory formalize this phenomenon and study it methodically?

  Research has shown that such parameters exist for many, previously classified as intractable, problems and when restricted to the case where they are bounded, there exist algorithms that justify their placement into the sphere of tractability. The related area, which has gained much attention recently, is called *Parameterized Complexity* and the algorithms designed in this setting are called *parameterized* (or *multivariate*) *algorithms*.
  As Parameterized Complexity is closely related to the topic discussed in this thesis, we will give some formal definitions and references in following subsection.

All the previously mentioned methods have been studied extensively in the last decades and each one of them constitutes a wide research area in the frame of Theoretical Computer Science. Of course ideas and techniques from any of these areas "flow" between them and researches are always interested in combining notions from some of or all the fields, as, for example, indicated (already in the title) by [28].

3.2 Parameterized Complexity and Algorithms

Let Σ be an alphabet (for example Σ can be the set $\{0, 1\}$) and let Σ^* (the *Kleene star* of Σ) be the set of all finite sequences with elements from Σ.

Parameterized languages and problems. We will call every subset L of $\Sigma^* \times \mathbb{N}$ a *parameterized language* and for every element $(x, k) \in L \subseteq \Sigma^* \times \mathbb{N}$ we will say that k is the parameter and x the *main input*. For every $k \in \mathbb{N}$, we call $L_k = \{(x, k) : (x, k) \in L\}$ the kth *slice* of L.
A decision problem Π is called a *parameterized problem* if any instance of it is encoded as a pair $(x, k) \subseteq \Sigma^* \times k$. We will say that (x, k) is a YES-instance for Π if (x, k) encodes an instance for which the question imposed in problem Π is answered positively, and will write $(x, k) \in \Pi$. Otherwise we will say that (x, k) is a NO-instance for Π if (x, k) and will write $(x, k) \notin \Pi$. If Π is a parameterized problem then it naturally defines the parameterized language

$$L_\Pi = \big\{(x, k) \in \Sigma^* \times k \mid (x, k) \in \Pi\big\}$$

Fixed-parameter tractability. We say that a parameterized problem Π is *fixed parameter tractable* if and only if there exists an algorithm $\mathcal{A}$ (or more formally a deterministic Turing Machine), a constant c, and a computable function f such that, for all $(x, k) \in \Sigma \times k$, $\mathcal{A}\big((x, k)\big)$ runs in time at most $f(k)|x|^c$ (where $|x|$ is the length of x) and

$$(x, k) \in L_\Pi \iff \mathcal{A}\big((x, k)\big) = 1$$

The class of all fixed-parameter tractable problems is called class FPT and is considered as the class of efficiently solvable problems in the world of Parameterized Complexity (can be thought of as the analog of P in terms of classical complexity).

There are also complexity classes that represent the intractable parameterized problems (such as the W-hierarchy, the A-hierarchy and XP) but their study is out of the scope of this thesis. For a nice introduction see [29] and for more advanced topics see [11], [10], and [6].

CHAPTER 4

DEFINITIONS

4.1 Plane graphs and pd-graphs

All graphs we consider are loop-less and may have multiple edges. We say that a graph is *plane* when it is embedded, without crossings between its edges, in the sphere $\mathbb{S}^2 = \{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 = 1\}$, and we treat it as its embedding in $\mathbb{S}^2$ or as the union of points that correspond to its vertices and edges. We will use the term *general graph* to stress that a graph is treated as a combinatorial structure and not as a topological (embedded) one. Let A be a subset of $\mathbb{R}^2$ (when we write $\mathbb{R}$ we refer to the plane with the standard euclidean metric). We define **int**(A) to be the interior of A, **cl**(A) its closure and **bnd**$(A) =$ **cl**$(A) \setminus$ **int**(A) its boundary.

We define a closed arc (resp. open arc) α to be a subset of $\mathbb{R}^2$ that is homeomorphic to $[0, 1]$ (resp. $(0, 1)$), meaning that there exists a bicontinuous function $f : [0, 1] \to \alpha$ (resp. $f : (0, 1) \to \alpha$). We call $f(0)$ and $f(1)$ the endpoints of arc α. For a closed arc α (that corresponds to function $f : [0, 1] \to \alpha$) we define **trim**(α) as the corresponding open arc $\alpha' = \alpha \setminus \{f(0), f(1)\}$. Observe that the faces of a plane graph are open sets and its edges are (open) arcs.

We say that two paths, P_1 and P_2, of a graph G are *disjoint* if none of the internal vertices of the one is a vertex of the other. A path is *non-trivial* if it contains at least two vertices. Given a graph G we denote by $\mathcal{P}(G)$ the set of all the paths in G.

Topological isomorphism of plane graphs. Let G be a graph and Γ be a plane graph. We denote by $\mathcal{C}(G)$ the set of the connected components of G. For every $f \in F(\Gamma)$ we denote by $B_\Gamma(f)$ the graph induced by the vertices and edges of Γ whose embeddings are subsets of **bd**(f).

We define a *closed walk* of a graph G to be a cyclic ordering $w = (v_1, \ldots, v_l, v_1)$ of vertices of $V(G)$ such that for any two consecutive vertices, say v_i, v_{i+1}, there is an edge between them in G, i.e., $\{v_i, v_{i+1}\} \in E(G)$. Note here that there may exist two distinct indices i, j such that $v_i, v_j \in w$ and $v_i = v_j$ (the walk can revisit a vertex). We say that a walk w of a plane graph Γ is *facial* if there exists $f_i \in F(\Gamma)$ and $\Theta_j \in \mathcal{C}(B_\Gamma(f_i))$ such that the vertices of w are the vertices of $V(\Theta_j)$ and the cyclic ordering of w indicates the way these vertices are met when making a closed walk along Θ_j while always keeping f_i on the same side of the walk.

Given that Γ is a plane graph and $\mathbf{w} = \{w_1, \dots, w_p\}$ is a non-empty set of closed walks of Γ, we say that $\mathbf{w}$ is a *facial mapping* if there exists some face f of Γ such that $\mathcal{C}(B_\Gamma(f)) = \{\Theta_1, \Theta_2, \dots, \Theta_p\}$ and w_j is a facial walk of Θ_j, $j \in [p]$. Given a plane graph Γ and $f \in F(\Gamma)$, we define $\mathbf{w}(f)$ as the facial mapping of Γ corresponding to f. Observe that for every face $f \in \Gamma(F)$, its facial mapping $\mathbf{w}(f)$ is unique (up to permutations).

Let Γ and Δ be two plane graphs. We say that Γ and Δ are *topologically isomorphic* if they are isomorphic via a bijection $g : V(\Gamma) \to V(\Delta)$ and there exists a function $h : F(\Gamma) \to F(\Delta)$, such that for every $f \in F(\Gamma)$, $g(\mathbf{w}(f)) = \mathbf{w}(h(f))$ (where $g(\mathbf{w}(f))$ is the result of applying g to every vertex of every closed walk in $\mathbf{w}$). We call the function $\alpha : V(\Gamma) \cup F(\Gamma) \to V(\Delta) \cup F(\Delta)$ such that $\alpha = g \cup h$, a *topological isomorphism* between Γ and Δ. Given two plane graphs G_1 and G_2 we say that they are the *same* graph if they are topologically isomorphic (and not just isomorphic).

Vertex dissolution. Let G be a graph and let $v \in V(G)$ such that $\deg_G(v) = 2$ and $N_G(v) = \{x, y\}$. We say that G' is the graph obtained from G after *dissolving* vertex v, if $V(G') = V(G) \setminus \{v\}$ and $E(G') = (E(G) \cup \{x, y\}) \setminus (\{v, x\} \cup \{v, y\})$. If G is plane and we dissolve a vertex $v \in V(G)$, we can just remove the point that represents v in the embedding of G and join the two arcs that correspond to the edges adjacent to v.

Primal-dual graphs. Let G be a plane graph. Observe that G does not necessarily have a unique embedding. We denote by $\mathbf{duals}(G)$ the set of all different duals of G (by "different'' we mean mutually not topologically isomorphic). Given a plane graph $J \in \mathbf{duals}(G)$, we define $\mathbf{pd}(G, J)$ as the plane graph obtained if we consider both embeddings of G and J such that in the resulting embedding, every edge e of $E(G)$ intersects its dual edge e^* in $E(J)$ at only one point. For each such intersection we introduce a new vertex, v_e, and we embed it on the intersection point. We call $\mathbf{pd}(G, J)$ a *primal-dual graph* or of G. The vertex set of $\mathbf{pd}(G, J)$ is naturally partitioned to the primal vertices, i.e., the vertices of G, the dual vertices, i.e., the vertices of J, and the crossing vertices in the set $\{v_e \mid e \in E(G)\}$. We denote these three sets by $V_p(\mathbf{pd}(G, J))$, $V_d(\mathbf{pd}(G, J))$, and $V_c(\mathbf{pd}(G, J))$ respectively and we say that two vertices of G are of the same *type* if they belong to the same set of the partition $\{V_p(\mathbf{pd}(G, J)), V_d(\mathbf{pd}(G, J)), V_c(\mathbf{pd}(G, J))\}$. For an example see Figure 4.1.

Given a path P of $\mathbf{pd}(G, J)$, we call it a *primal path* (resp. *dual path*) if $V(P) \subseteq V_p(\mathbf{pd}(G, J)) \cup V_c(\mathbf{pd}(G, J))$ (resp. $V(P) \subseteq V_d(\mathbf{pd}(G, J)) \cup V_c(\mathbf{pd}(G, J))$) and not both its endpoints are in $V_c(\mathbf{pd}(G, J))$. Notice that a primal path in $\mathbf{pd}(G, J)$ may intersect with a dual path only in vertices in $V_c(\mathbf{pd}(G, J))$ and that the vertices of a primal (resp. dual) path alternate between primal (resp. dual) type and crossing type. Given a cycle C in **pd**(G, J) we say that C is of *primal type* or is a *primal cycle* (resp. *dual type* or is a *dual cycle*) if C is the union of two primal (resp. dual) paths of G.

We also define the *primal edges* (resp. *dual edges*) of **pd**(G, J) to be the set $E_p(\mathbf{pd}(G, J)) = \{\{u, v\} \in E(\mathbf{pd}(G, J)) \mid \{u, v\} \cap V_d(\mathbf{pd}(G, J)) = \emptyset\}$ (resp. $E_d(\mathbf{pd}(G, J)) = \{\{u, v\} \in E(\mathbf{pd}(G, J)) \mid \{u, v\} \cap V_p(\mathbf{pd}(G, J)) = \emptyset\}$). Clearly, this is a partition of $E(\mathbf{pd}(G, J))$, as any edge has exactly one endpoint in $V_c(\mathbf{pd}(G, J))$. Given a two vertices $s, t \in V(\mathbf{pd}(G, J))$, we say that (s, t) is a *primal* (resp. *dual*) *pair* of **pd**(G, J) if $|\{s, t\} \cap V_c(\mathbf{pd}(G, J))| \leq 1$ and $\{s, t\} \cap V_d(\mathbf{pd}(G, J)) = \emptyset$ (resp. $\{s, t\} \cap V_p(\mathbf{pd}(G, J)) = \emptyset$).

A plane graph is called *pd-graph* if it is the primal-dual graph **pd**(G, J) for some plane graph G and some $J \in \mathbf{duals}(G)$. Notice that if G is a connected plane graph then **duals**(G) contains only one graph which we denote by G^* and therefore its pd-graph is also unique, the primal-dual graph **pd**(G, G^*), and we denote it by **pd**(G).

When we are given a pd-graph we will assume that the partition of its vertices to primal, dual, and crossing is also given and we will denote by G^p (resp. G^d) the graph obtained if we dissolve all crossing vertices in $G[V_p(G)]$ (resp. $G[V_d(G)]$). The plane graph G^p (resp. G^d) is the *primal part* (resp. *dual part*) of G.

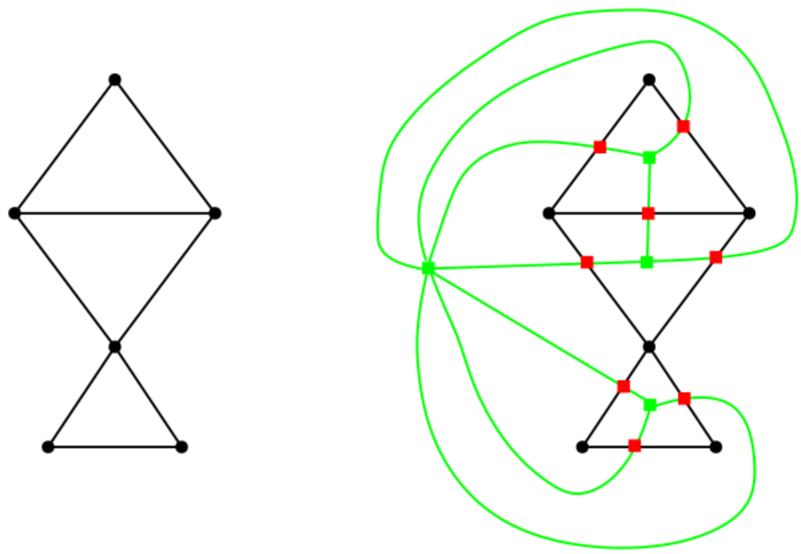

Figure 4.1: A plane (connected) graph on the left and the corresponding (unique) pd-graph on the right. The black dots (resp. lines) represent the primal vertices (resp. edges), the red squares the crossing vertices and the green squares (resp. lines) the dual vertices (resp. edges).

We are now in the position to formally define the problem of interest:

pd-Disjoint Paths (pd-DP)
Input: A pd-graph $G = (V, E)$, a collection $\mathcal{P} = \{(s_i, t_i) \in 2^{V_2}, i \in \{1, \ldots, k\}\}$ of k disjoint pairs of terminals of G and a mapping $\tau : \mathcal{P} \to \{p, d\}$.
Question: Are there k pairwise vertex-disjoint paths $P_1, \ldots, P_k$ in G such that for $i \in \{1, \ldots, k\}$, P_i has endpoints s_i and t_i and additionaly P_i is a primal path iff $\tau\big((s_i, t_i)\big) = p$?

Primal-dual contraction. Let G be a pd-graph. We call two vertices x_1 and x_2 *adjoined* if they have as a common neighbor a vertex $y \in V_c(G)$ and they either are both in $V_p(G)$ or are both in $V_d(G)$. We then say that vertices x_1 and x_2 are *adjoined through* vertex y. The operation of *primal-dual contraction* of two adjoined vertices x_1 and x_2 is defined as follows:

1. Delete vertices y, x_1, x_2 from G.

2. Add a new vertex $x_{1,2}$ and the edges $\{\{x_{1,2}, v\} \mid v \in (N_G(x_1) \cup N_G(x_2)) \setminus \{y\}\}$.

The operation of the *primal-dual removal* of x_1 and x_2 is the operation of the *primal-dual contraction* of the vertices in $N_G(y) \setminus \{x_1, x_2\}$.
Observe that the operation of the primal-dual contraction of the adjoined vertices x_1, x_2 in G, results to a new pd-graph, which we denote by $G^{(x_1,x_2)}$. For an example see Figure 4.1.

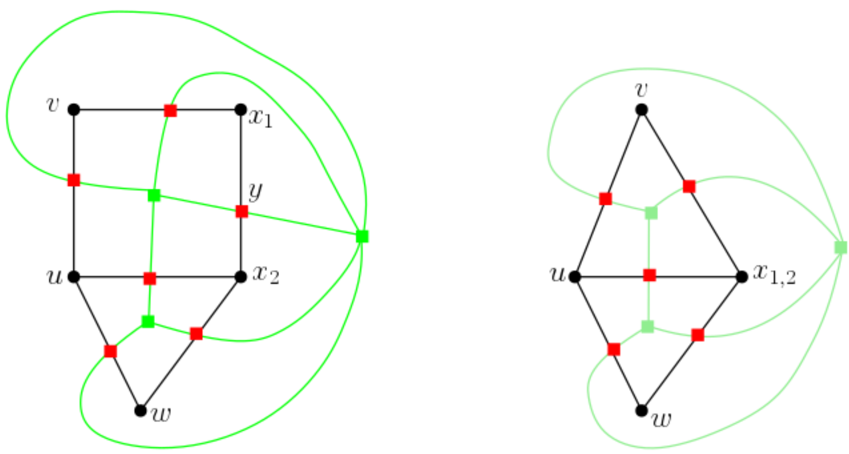

Figure 4.2: An illustration of the primal-dual contraction of two adjoined vertices x_1 and x_2. The pd-graph on the left is the one before the contraction and the one on the right depicts the pd-graph obtained after the contraction.

Alternating cycles in pd-graphs. Let G be a pd-graph and let $\mathcal{C} = \{C_0, \ldots, C_r\}$ be a sequence of cycles in G. We say that the sequence of cycles $\mathcal{C}$ is *alternating*, if the cycles in $\mathcal{C}$ have alternating types, e.g. C_i is a primal cycle iff

$i \equiv 0 \pmod 2$. If C_i is primal (resp. dual) we also refer to it as C_i^p (resp. C_i^d) to stress its type. We call $\{C \in \mathcal{C} \mid C = C_i^p$ for some $i \in \{0, \dots k\}\}$ and $\{C \in \mathcal{C} \mid C = C_i^d$ for some $i \in \{0, \dots k\}\}$ the *primal subsequence* and the *dual subsequence* of $\mathcal{C}$, respectively.

Tight concentric cycles. Let G be a plane graph and let D be a disk that is the closed interior of some cycle C of G. We say that D is *internally chordless* if there is no path in G whose endpoints are vertices of C and all the others belong in **int**(D).

Let $\mathcal{C} = \{C_0, \dots, C_r\}$, be a sequence of cycles in G. We denote by D_i the closed interior of $C_i, i \in \{0, \dots, r\}$, and we say that $\mathcal{D} = \{D_0, \dots, D_r\}$ is the *disc sequence* of $\mathcal{C}$. We call $\mathcal{C}$ *concentric*, if for all $i \in \{0, \dots, r-1\}$, the cycle C_i is contained in the open interior of D_{i+1}. We also define $\Delta(\mathcal{C})$ as the intersection of the closed interior of C_1 and the closed exterior of C_r. The sequence $\mathcal{C}$ of concentric cycles is *tight* in G, if, in addition,

- D_0 is *internally chordless*,
- for every $i \in \{0, \dots, r-1\}$, there is no cycle of G that is contained in $D_{i+1} \setminus D_i$ and whose closed interior D has the property $D_i \subset D \subset D_{i+1}$.

Let H be a pd-graph. We say that a sequence $\mathcal{A}$ of concentric cycles of the same type in H is *primal-tight* (resp. *dual-tight*) if $\mathcal{A}$ is tight in $H \setminus V_d(H)$ (resp. in $H \setminus V_p(H)$).

Let $\mathcal{C} = \{C_0, \dots, C_r\}$ be a sequence of primal-tight concentric cycles in H and for every $i \in \{0, \dots r-1\}$ there exists a dual cycle C_i^* in H such that $C_i^* \subset D_{i+1} \setminus D_i$. We say that the sequence $\mathcal{C}^* = \{C_0^*, \dots, C_{r-1}^*\}$ is *tight with respect to* $\mathcal{C}$, if for every $i \in \{0, \dots r-1\}$ there does not exist a cycle C of dual type in H such that $D_i^* \subset D \subset D_i$. We define the tightness of a sequence of primal cycles with respect to a sequence of dual-tight concentric cycles in a symmetric way. We say that an alternating sequence $\mathcal{Z} = \{Z_1^p, Z_2^d, \dots, Z_{r-1}^p, Z_r^d\}$ (if $\mathcal{Z}$ starts with a cycle of dual type we again have a symmetric definition) of concentric cycles in H is *tight* in H, if $\mathcal{Z}^p = \{Z_1^p, Z_3^p \dots, Z_{r-1}^p\}$ is primal-tight in H and $\mathcal{Z}^d = \{Z_2^d, Z_4^d \dots, Z_r^d\}$ is tight with respect to $\mathcal{Z}^p$.

It is not hard to verify that, if $\mathcal{C} = \{C_1^p, C_2^d, \dots, C_{r-1}^p, C_r^d\}$ is an alternating sequence of concentric cycles in H and $\mathcal{C}^p = \{C_1^p, C_3^p, \dots C_{r-1}^p\}$ is a primal-tight sequence of concentric cycles in H, then there exist dual cycles ${C'}_2^d, {C'}_4^d, \dots {C'}_r^d$ in H such that $\mathcal{C}_d = \{{C'}_2^d, {C'}_4^d, \dots {C'}_r^d\}$ is tight with respect to $\mathcal{C}^p$. We call $\mathcal{C}' = \{C_1^p, {C'}_2^d, \dots, C_{r-1}^p, {C'}_r^d\}$ an *alternating tight sequence of concentric cycles that corresponds to* $\mathcal{C}$.

4.2 Linkages in pd-graphs

Linkages and pd-linkages. A *linkage* in a graph G is a subgraph L of G whose connected components are all paths. The *terminals* of a linkage L are the endpoints of the paths in $\mathcal{P}(L)$, and the *pattern* of L is the set of all tuples (s, t), where s and t are distinct terminals such that there exists a path with endpoints s and t in G (resp. in H). We say that two linkages are *equivalent* if they have the same pattern.

Let H be a pd-graph. A *primal-dual linkage* or *pd-linkage* L of H is a collection of pairwise vertex disjoint, non-trivial paths of H, each of which is either primal or dual. The definition of the *paths* and the *terminals* of L are the same as in the case of a L being a linkage. The *pattern* of L is the set of all triples $(s, t, \tau_{s,t})$, where s and t are distinct terminals such that there exists a path with endpoints s and t in L and $\tau_{s,t}$ is the type of this path. We say that two pd-linkages are *equivalent* if they have the same pattern. For a visualization of these notions see Figure 4.2.

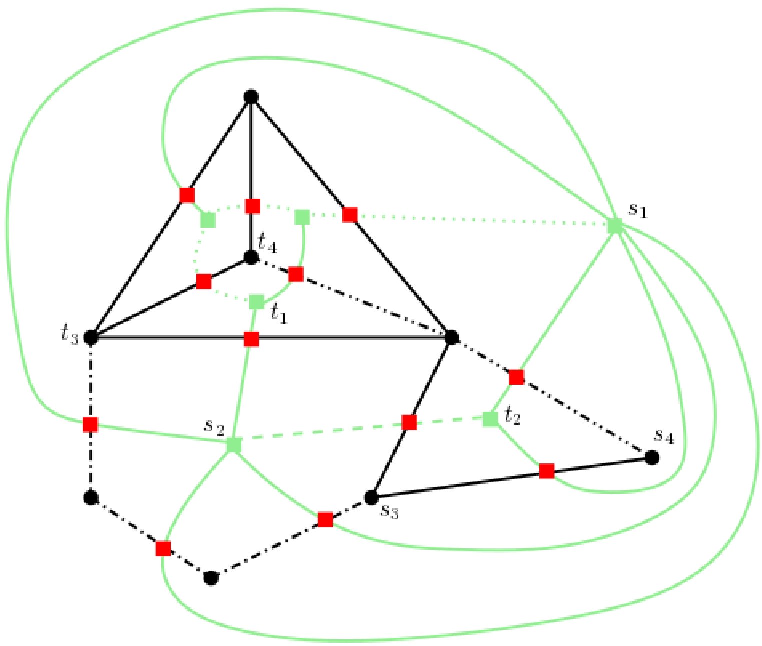

Figure 4.3: A pd-linkage L of a pd-graph G, with pattern $\{(s_i, t_i) : i \in [4]\}$ and order 4. The terminals of L are $s_1, t_1, s_2, t_2, s_3, t_3, s_4$ and t_4. The dashed and dotted paths indicate the paths of the pd-linkage. Different lining corresponds to different path, while the color black corresponds to primal paths and the color green to dual paths of L.

Segments. Let G be a pd-graph and let C be a cycle in G whose closed interior is D. Given a path P in G we say that a subpath P_0 of P is a D-segment of P, if P_0 is a non-empty (possibly edgeless) path obtained by intersecting P with D. For a pd-linkage L of G we say that a path P_0 is a D*-segment* of L, if P_0 is a D*-segment* of some path P in $\mathcal{P}(L)$.

CL**-configurations.** Given a pd-graph G, we say that a pair $\mathcal{Q} = (\mathcal{C}, L)$ is a *CL-configuration* of G of *depth* r if $\mathcal{C} = \{C_0, \ldots, C_r\}$ is an alternating sequence of concentric cycles in G, L is a pd-linkage of G, and the intersection of the set of terminals of L with D_r is empty. A *segment* of $\mathcal{Q}$ is any D_r-segment of L. The *eccentricity* of a segment P of $\mathcal{Q}$ is the minimum i such that $V(C_i \cap P) \neq \emptyset$. A segment of $\mathcal{Q}$ is *extremal* if it is has eccentricity r, i.e., it is a subgraph of C_r. Given a cycle $C_i \in \mathcal{C}$ and a segment P of $\mathcal{Q}$ we define the *i-chords* of P as the connected components of $P \cap \mathbf{int}(D_i)$ (notice that i-chords are open sets). For every i-chord X of P, we define the *i-semichords* of P as the connected components of the set $X \setminus D_{i-1}$ (notice that i-semichords are open sets). Given a segment P that does not have any 0-chord, we define its *zone* as the connected component of $D_r \setminus P$ that does not contain the open interior of D_0 (a zone is an open set). A CL-configuration $\mathcal{Q} = (\mathcal{C}, L)$ is called *reduced* if the graph $L \cap \bigcup\mathcal{C}$ is edgeless.

Observe that if a primal path $P \in \mathcal{P}(L)$ (the argument is the same for a dual path) has common edges with some cycle, say $C_i \in \mathcal{C}$, then C_i is a primal cycle (as P cannot contain any vertices in $V_d(G)$) and every connected component $R = C_i \cap P$ contains an even number of edges. That is because R is a path with both endpoints, u, v, in $V_p(G)$, as the neighbors of u and v in $V(P) \setminus V(R)$ are in $V_c(G)$ (because they are both vertices in the intersection of the primal path P and some dual cycle).

Let $\mathcal{Q} = (\mathcal{C}, L)$ be a CL-configuration of G and let $E^\bullet$ be the set of all edges of the graph $L \cap \bigcup\mathcal{C}$ (remember that each edge in the pd-graph G has exactly one endpoint in $V_c(G)$). We then define G^* as the graph obtained if we apply primal-dual contractions (or removals) to all vertices in $V(\bigcup\mathcal{C})$ that are adjoined through a vertex $c \in V_c$ that is an endpoint of some edge in $E^\bullet$. The previous observation explains why this can be done The sequence of cycles $\mathcal{C}$ is modified into a sequence of cycles in G^*, say $\mathcal{C}^*$ and L is modified into a pd-linkage in G^*, say L^*. Notice that $\mathcal{Q}^* = (\mathcal{C}^*, L^*)$ is a reduced CL-configuration of G^*. We call $(\mathcal{Q}^*, G^*)$ the *reduced pair* of G and $\mathcal{Q}$.

Cheap pd-linkages. Let G be a pd-graph and $\mathcal{Q} = (\mathcal{C}, L)$ be a CL-configuration of G of depth r. We define the function $c : \{L \mid L \text{ is a linkage of } G\} \to \mathbb{N}$ so that

$$c(L) = |E(L) \setminus \bigcup_{i \in \{0,\ldots,r\}} E(C_i)|.$$

A pd-linkage L of G is $\mathcal{C}$-cheap, if there is no other CL-configuration $\mathcal{Q}' = (\mathcal{C}, L')$ such that $c(L) > c(L')$. Intuitively, the function c penalizes every edge of the linkage that does not lie on some cycle C_i.

Observation 4.2.1. *Let $\mathcal{Q} = (\mathcal{C}, L)$ be a CL-configuration of some pd-graph G and let $\big(G^* \mathcal{Q}^* = (\mathcal{C}^*, L^*)\big)$ be the reduced pair of G and $\mathcal{Q}$. Then*

- *If L is $\mathcal{C}$-cheap, then L^* is $\mathcal{C}^*$-cheap.*
- *If $\mathcal{C}$ is tight in G, then if $\mathcal{C}^*$ is tight in G^*.*

Convex CL-configurations. Let G be a pd-graph and let $\mathcal{Q} = (\mathcal{C}, L)$ be a CL-configuration of G. A segment P of $\mathcal{Q}$ is *convex* if the following three conditions are satisfied:

(i) it has no 0-chords and

(ii) for every $i \in \{1, \ldots, r\}$, the followings hold:

a. P has at most one i-chord

b. if P has an i-chord, then $P \cap C_{i-2} \neq \emptyset$.

c. Each i-chord of P has exactly two i-semichords.

(iii) If P has eccentricity $i < r - 1$, there is some other segment of $\mathcal{Q}$ with eccentricity $i + 2$ inside the zone of P.

We say $\mathcal{Q}$ is *convex* if all its segments are convex.

Observation 4.2.2. *Let $\mathcal{Q} = (\mathcal{C}, L)$ be a CL-configuration and let $(G^*, \mathcal{Q}^* = (\mathcal{C}^*, L^*))$ be the reduced pair of G and $\mathcal{Q}$. Then $\mathcal{Q}$ is convex if and only $\mathcal{Q}^*$ is convex.*

We will need the following topological lemma.

Lemma 4.2.1 (Lemma 2 of [2])**.** *Let Δ_1, Δ_2 be closed disks of $\mathbb{R}^2$ where* **int**$(\Delta_1) \cap$ **int**$(\Delta_2) = \emptyset$ *and such that $\Delta_1 \cup \Delta_2$ is also a closed disk. Let $\Delta_3 = \mathbb{R}^2 \setminus$* **int**$(\Delta_1 \cup \Delta_2)$ *and let $Y =$* **bnd**$(\Delta_3) \cap \Delta_2$ *and $Q =$* **trim**$(\Delta_1 \cap \Delta_2)$*. P be a closed arc of $\mathbb{R}^2$ whose endpoints are not in $\Delta_1 \cup \Delta_2$ and such that $Y \cap P = \emptyset$ and $Q \cap P \neq \emptyset$. Then* **int**$(\Delta_1) \cap P$ *has at least two connected components.*

Proof. Let q be some point in $Q \cap P$. Let Q' be an open arc that is a subset of **int**(Δ_1) and has the same endpoints as Y. Notice that q and x belong to different open disks defined by the cycle $Q' \cup Y$. Therefore P should intersect Q' or Y. As $Y \cap P = \emptyset$, P intersects Q'. As $Q' \subseteq$ **int**(Δ_1), **int**$(\Delta_1) \cap P$ has at least one connected component.

Assume now that **int**$(\Delta_1) \cap P$ has exactly one connected component. Clearly, this connected component will be an open arc I such that at least one of the endpoints of I, say q, belongs to Q. Moreover, there is a subset P' of P that is a closed arc where $P' \cap I = \emptyset$ and whose endpoints are q and one of x and y, say y. As **int**$(\Delta_1) \cap P$ has exactly one connected component, it holds that $P' \cap$ **int**$(\Delta_1) = \emptyset$. Let Q' be an open arc that is a subset of **int**(Δ_1) and has the same endpoints as Y. Notice that q and y belong to different open disks defined by the cycle $Q' \cup Y$. Therefore P' should intersect **int**(Δ_1) or Y, a contradiction as $P' \subseteq P$ and $Y \cap P = \emptyset$. □

Lemma 4.2.2. *Let G be a pd-graph and let $\mathcal{Q} = (\mathcal{C}, L)$ be a CL-configuration of G where $\mathcal{C}$ is tight in G and L is $\mathcal{C}$-cheap. Then $\mathcal{Q}$ is convex.*

Proof. Let $\mathcal{C} = \{C_1, \ldots, C_r\}$. It suffices to prove that any primal segment is convex, as the argument is identical for the dual case. Moreover, by 4.2.1 and 4.2.2 we may assume that $\mathcal{Q}$ is reduced.
Condition (i) follows from the tightness of $\mathcal{C}$ (a 0-chord implies the existence of a path in D_0 with endpoints in C_0). Similarly for condition (ii).b, as the existence of an i-chord that does not intersect C_{i-2} contradicts the tightness of $\mathcal{C}$ (actually contradicts the primal-tightness of the primal subsequence of $\mathcal{C}$). Moreover,

condition (iii) follows from the hypothesis that L is $\mathcal{C}$-cheap: Let P be a primal segment with eccentricity i. If there does not any segment with eccentricity $i+2$ then we contradict the cheapness of L by using C_{i+2} to reroute P, obtaining an equivalent linkage.

Let $i \in \{0, \ldots, r\}$ be the least index such that one of conditions (ii).a and (ii).c. is violated and let W be a primal segment of $\mathcal{Q}$ containing an i-chord X that violates one of the conditions. We distinguish two cases:

Case 1: Condition (ii).c is violated. From condition (ii).b, $X \setminus D_{i-1}$ contains at least three i-semichords of X. Let J_1 be the biconnected outerplanar graph defined by the union of C_{i-1} and those i-semichords of X that do not intersect C_i. By the minimality of i, J_1 has at least one internal edge and therefore at least two simplicial faces, and there exist exactly two i-semichords of X, say K_1 and K_2, that intersect C_i, which belong to the same face, say F_1, of J_1. We define Δ_2 to be the closure of a simplicial face of J_1 that is different from F_1.

Case 2: Condition (ii).c holds while condition (ii).a is violated. Let J_2 be the biconnected outerplanar graph defined by the union of C_{i-1} and the connected components of $W \setminus D_{i-1}$ that do not intersect C_r. Notice that the remaining connected components of $W \setminus D_{i-1}$ are exactly two, say K_1 and K_2, and are subsets of the same face, say F_1, of J_2. Moreover, as there exist at least two i-chords in W, J_2 contains at least one internal edge and therefore at least two simplicial faces. We define Δ_2 to be the closure of a simplicial face of J_2 that is different from F_1.

The remaining part of the proof works for both of the above cases: We set $\Delta_1 = D_{i-1}$, $\Delta_3 = \mathbb{R}^2 \setminus \textbf{int}(\Delta_1 \cup \Delta_2)$, $Y = \textbf{bnd}(\Delta_3) \cap \Delta_2$, and $Q = \textbf{trim}(\Delta_1 \cap \Delta_2)$. It is clear from the definition of Y that it is a subpath of W, thus $Y \subseteq W$.

Suppose that $L \cap Q = \emptyset$. Then consider W' as a path in $W \cup Q$ such that $Q \subseteq W'$ and W' has the same endpoints with W, and define $L' = (L \setminus W) \cup W'$. Linkage L' is equivalent to L (they clearly have the same pattern) and $c(L') < c(L)$ (as we got rid of at least one edge L that was not an edge of a cycle in $\mathcal{C}$), a contradiction to L being $\mathcal{C}$-cheap. Thus, $L \cap Q \neq \emptyset$ meaning that L contains a segment P for which $P \cap Q \neq \emptyset$. We again distinguish two cases:

Case a: $W \neq P$. Then, $W \cap P = \emptyset$ (the paths of a linkage are disjoint) and as $Y \subseteq W$, $Y \cap P = \emptyset$. Clearly, P is a path whose endpoints are not in $\Delta_1 \cup \Delta_2$ and Lemma 4.2.1 can be applied. Thus, $P \cap \textbf{int}(\Delta_1)$ has at least two connected components, therefore P has at least two $(i-1)$-chords. If $i = 1$ then P has a 0-cord which violates condition (i) (which should hold for every segment), and if $i > 1$ condition (ii).a is violated, which contradicts the minimality of i.

Case b: $W = P$. Let p_1 and p_2 be the endpoints of $\textbf{cl}(Q)$. Then, there exists two disjoint closed arcs Z_1 and Z_2 with endpoints p_1, p_1' and p_2, p_2', respectively, such that

- $Z_i \subseteq \textbf{cl}(Q), i \in \{1, 2\}$, and
- $P \cap Z_i = \{p_i\}, i \in \{1, 2\}$.

Consider also a closed arc Y' that is a subset of $\textbf{int}(\Delta_2) \cup \{p_1', p_2'\}$ that does not intersect L and whose endpoints are p_1' and p_2'. Let now $\Delta_1' = \Delta_1$, Δ_2' be the closed disk defined by the cycle **cl**$(Q \setminus (Z_1 \cup Z_2)) \cup Y'$ that is a subset of Δ_2. Let also $\Delta'_3 = \mathbb{R}^2 \setminus$ **int**$(\Delta'_1 \cup \Delta'_2)$ and $Q' =$ **trim**$(\Delta'_1 \cap \Delta'_2)$. As Y' does not intersect L, we obtain $Y' \cap P = \emptyset$. Observe that Z_1, Q', Z_2 form a partition of Q. As $Q \cap P \neq \emptyset$ and $(Z_i \setminus \{p_i\}) \cap P = \emptyset, i \in \{1, 2\}$, we conclude that $Q' \cap P \neq \emptyset$.

By applying Lemma 4.2.1, **int**$(\Delta'_1) \cap P$ has at least two connected components. Therefore P has at least two $(i-1)$-chords and this yields a contradiction, as in Case a. □

CHAPTER 5

OUTSIDE OF THE OUTER CYCLE

5.1 Bounding the number of extremal segments

In this section, following section 3.3 of [2], we prove that the number of extremal segments of a CL-configuration $\mathcal{Q} = (\mathcal{C}, L)$ of a pd-graph, is at most a linear function of the order of the pd-linkage L. The arguments in this part are independent of G being a pd-graph or just a plane graph, thus we can argue as if G was plane and the pd-linkages are linkages of G.

Out-segments, hairs, and flying hairs. Let G be a pd-graph and $\mathcal{Q} = (\mathcal{C}, L)$ be a CL-configuration of G of depth r. An *out-segment* of L is a subpath P' of a path $P \in \mathcal{P}(L)$ such that the endpoints of P' are in C_r and the internal vertices of P' are not in D_r. A *hair* of L is a subpath P' of a path in $\mathcal{P}(L)$ such that one endpoint of P' is in C_r, the other is a terminal of L, and the internal vertices of P' are not in D_r. A *flying hair* of L is a path in $\mathcal{P}(L)$ that does not intersect C_r.

Given a pd-linkage L of G and a closed disk D of $\mathbb{R}^2$, we define $\textbf{out}_D(L)$ to be the graph obtained from $(L \cup \textbf{bnd}(D)) \setminus \textbf{int}(D)$ after dissolving all vertices of degree two. For example $\textbf{out}_{D_r}(L)$ is a plane graph consisting of the out-segments, the hairs, and the flying hairs of L and what remains from C_r after dissolving its vertices that do not belong in L. For an example see Figure 5.1.

Let f be a face of $\textbf{out}_{D_r}(L)$ that is different from $\textbf{int}(D_r)$. We say that f is a *cave of* $\textbf{out}_{D_r}(L)$ if the union of the out-segments and extremal segments in the boundary of f form a connected set. Recall that a segment of $\mathcal{Q}$ is extremal if it is has eccentricity r, i.e., it is a subpath of C_r.

Given a plane graph G, we say that two edges e_1 and e_1 are *cyclically adjacent* if they have a common endpoint x and appear consecutively in the cyclic ordering of the edges incident to x, as defined by the embedding of G. A subset E of $E(G)$ is *cyclically connected* if for every two edges e and e' in E there

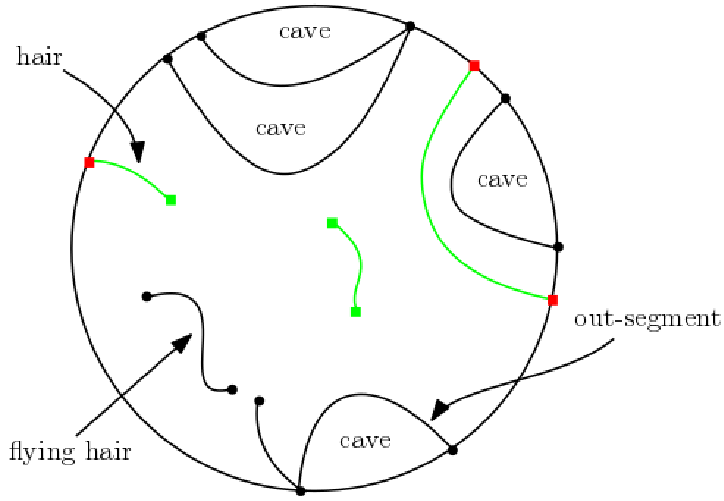

Figure 5.1: The interior of the black (primal) cycle corresponds to $\mathbf{out}_D(L)$. The black dots (edges) correspond to primal vertices (paths), the green squares (edges) to dual vertices (paths), and the re squares to crossing vertices.

exists a sequence of edges $e_1, \ldots, e_r \in E$ where $e_1 = e$, $e_r = e'$ and for each $i \in \{1, \ldots, r-1\}$ e_i and e_{i+1} are cyclically adjacent.

Let $\mathcal{Q} = (\mathcal{C}, L)$ be a CL-configuration. We say that $\mathcal{Q}$ is *touch-free* if for every path P in $\mathcal{P}(L)$, the number of the connected components of $P \cap C_r$ is not 1.

Lemma 5.1.1. *Let G be a pd-graph and $\mathcal{Q} = (\mathcal{C}, L)$ be a touch-free CL-configuration of G where $\mathcal{C}$ is tight in G and L is $\mathcal{C}$-cheap. The number of extremal segments of $\mathcal{Q}$ is at most $2 \cdot |\mathcal{P}(L)| - 2$.*

Proof. The proof of this Lemma is similar to the one of Lemma 4 in [2] but we include it for this sake of completeness.
Let $(G^*, \mathcal{Q}^* = (\mathcal{C}^*, L^*))$ be the reduced pair of G and $\mathcal{Q}$. Notice that, by Observation 4.2.1, $\mathcal{C}^*$ is tight in G and L^* is $\mathcal{C}^*$-cheap. Moreover, it is easy to see that $\mathcal{Q}^*$ is touch-free and $\mathcal{Q}$ and $\mathcal{Q}^*$ have the same number of extremal segments which are just vertices of C_r in $\mathcal{Q}^*$. Therefore, it is sufficient to prove that the lemma holds for $\mathcal{Q}^*$. Let ρ be the number of extremal segments of $\mathcal{Q}^*$.

Let $J = \mathbf{out}_{D_r^*}(L^*)$ and $k = |\mathcal{P}(L^*)|$. Notice that the number of extremal segments of $\mathcal{Q}^*$ is equal to the number of vertices of degree 4 in J. The terminals of L^* are partitioned in three families

- *flying* terminals, T_0: endpoints of flying hairs.
- *invading* terminals T_1: these are endpoints of hairs whose non terminal endpoint has degree 3 in J
- *bouncing* terminals T_2: these are endpoints of hairs whose non terminal endpoint has degree 4 in J.

A hair containing an invading and bouncing terminal is called *invading* and *bouncing hair* respectively.

Recall that $|T_0| + |T_1| + |T_2| = 2k.$

Claim 1. The number of caves of J is at most the number of invading terminals.

Proof of claim 1. Clearly, a hair cannot be in the common boundary of two caves as it can be in the interior of a unique face of J. Therefore it is enough to prove that the boundary of every cave f contains at least one invading hair. To see this, consider the open arc R obtained if we remove from $\textbf{bnd}(f)$ all the points that belong to out-segments. Clearly, R results from a subpath R^+ of C_r^* after removing its endpoints, i.e., $R = \textbf{trim}(R^+)$.

Notice that because f is a cave, R is a non-empty connected subset of C_r^*. Moreover, $R \cap L^*$ is non-empty, otherwise $L^{*\prime} = (L^* \setminus (\textbf{bnd}(f)) \cup R$ is also a linkage with the same pattern as L^* where $c(L^{*\prime}) < c(L^*)$, a contradiction to the fact that L^* is $\mathcal{C}^*$-cheap. Let Y be a connected component of $R \cap L^*$. As $\mathcal{Q}^*$ is reduced, Y consists of a single vertex in the open set R. Notice that Y is a subpath of a segment Y' of $\mathcal{Q}^*$. We claim that Y' is not extremal. Suppose to the contrary that Y' is extremal. Then $Y' = Y$ and Y is a subset of the union of all extremal segments and out-segments in the boundary of f. This contradicts the fact that $Y \subseteq R$.

By Lemma 4.2.2, $\mathcal{Q}^*$ is convex, thus one of the endpoints of Y' is in Y and therefore in R as well. Let P be the path of L^* that contains Y'. Because all endpoints of paths in $\mathcal{P}(L^*)$ lie outside D_r^*, the set $P' = P \cap (\textbf{bnd}(f) \setminus D_r^*)$ is non empty and therefore, its closure P'^+ is either a hair or an out-segment of J. Assume that P'^+ is an out-segment. Then, again, Y is a subset of the union of all extremal segments and out-segments in the boundary of f and this contradicts the fact that $Y \subseteq R$. For the same reason, it cannot be a bouncing hair and therefore it is an invading hair. This completes the proof of Claim 1.

Let J^- be the graph obtained from J by removing all hairs and notice that J^- is a biconnected outerplanar graph. Let S be the set of vertices of J^- that have degree 4. Notice that, because $\mathcal{Q}^*$ is touch-free, $|S|$ is equal to the number of vertices of J that have degree 4 (which is equal to the number of extremal segments) minus the number of bouncing terminals, i.e., $|S| = \rho - |T_2|$. Therefore,

$$\rho = |T_2| + |S|. \tag{5.1}$$

Notice that if we remove from J^- all the edges of C_r^*, the resulting graph is a forest Ψ whose connected components are paths. Observe that none of these paths is a trivial path because $\mathcal{Q}^*$ is touch-free. We denote by $\kappa(\Psi)$ the number of connected components of Ψ. Let F be the set of faces of J^- that are different from D_r^*. F is partitioned into the faces that are caves, namely F_1 and the non-cave faces, namely F_0. By the Claim 1, $|F_1| \leq |T_1|$.

To complete the proof, it is enough to show that

$$|S| \quad \leq \quad |T_1| - 2 \tag{5.2}$$

Indeed the truth of (5.2) along with (5.1), would imply that ρ is at most $|T_2|+|S| \leq |T_2| + |T_1| - 2 \leq |T| - 2 = 2k - 2$.

We now return to the proof of (5.2). For this, we need two more claims.

Claim 2: $|F_0| \leq \kappa(\Psi) - 1$.
Proof. We use induction on $\kappa(\Psi)$. Let $K_1, \ldots, K_{\kappa(\Psi)}$ be the connected components of Ψ. If $\kappa(\Psi) = 1$ then all faces in F are caves, therefore $|F_0| = 0$ and we are done. Assume now that Ψ contains at least two connected components.

We assert that there exists at least one connected component K_h of Ψ with the property that only one non-cave face of J^- contains edges of K_h in its boundary. To see this, consider the weak dual T of J^-. Recall that, as J^- is biconnected, T is a tree. Let K_i^* be the subtree of T containing the duals of the edges in $E(K_i)$, $i \in \{1, \ldots, \kappa(\Psi)\}$, and observe that $E(K_1^*), \ldots, E(K_{\kappa(\Psi)}^*)$ is a partition of $E(T)$ into $\kappa(\Psi)$ cyclically connected sets. We say that a vertex of T is *rich* if it is incident with edges in more than one members of $\{K_1^*, \ldots, K_{\kappa(\Psi)}^*\}$, otherwise it is called *poor*. Notice that a vertex of T is rich if and only if its dual face in J^- is a non-cave. We call a subtree K_i^* *peripheral* if $V(K_i^*)$ contains at most one rich vertex of T. Notice that the claimed property for a component in $\{K_1, \ldots, K_{\kappa(\Psi)}\}$ is equivalent to the existence of a peripheral subtree in $\{K_1^*, \ldots, K_{\kappa(\Psi)}^*\}$. To prove that such a peripheral subtree exists, consider a path P in T intersecting the vertex sets of a maximum number of members of $\{K_1^*, \ldots, K_{\kappa(\Psi)}^*\}$. Let e^* be the first edge of P and let K_h^* be the unique subtree whose edge set contains e^*. Because of the maximality of the choice of P, $V(K_h^*)$ contains exactly one rich vertex v_h, therefore K_h^* is peripheral and the assertion follows. We denote by f_h the non-cave face of J^- that is the dual of v_h.

Let H^- be the outerplanar graph obtained from J^- after removing the edges of K_h. Notice that this removal results in the unification of all faces that are incident to the edges of K_h, including f_h, to a single face f^+. By the inductive hypothesis the number of non-cave faces of H^- is at most $\kappa(\Psi)-2$. Adding back the edges of K_h in J^- restores f_h as a distinct non-cave face of J^-. If f^+ was a non-cave of H^- then $|F_0|$ is equal to the number of non-cave faces of H^-, else $|F_0|$ is one more than this number. In any case, $|F_0| \leq \kappa(\Psi) - 1$, and the claim follows.

Claim 3: $|V(\Psi)| \leq |T_1| + 2 \cdot \kappa(\Psi) - 2$.
Proof. Let T be the weak dual of J^-. Observe that $|F_0|+|F_1| = |F| = |V(T)| = |E(T)| + 1 = |E(\Psi)| + 1 = |V(\Psi)| - \kappa(\Psi) + 1$. Therefore $|V(\Psi)| = |F_0| + |F_1| + \kappa(\Psi) - 1$. Recall that, by Claim 1, $|F_1| \leq |T_1|$ and, taking into account Claim 2, we conclude that $|V(\Psi)| \leq |T_1| + 2 \cdot \kappa(\Psi) - 2$. Claim 3 follows.

Notice now that a vertex of J^- has degree 4 iff it is an internal vertex of some path in Ψ. Therefore, as all connected components of Ψ are non-trivial paths, it holds that $|V(\Psi)| = |S| + |L(\Psi)| = |S| + 2 \cdot \kappa(\Psi)$, where $L(\Psi)$ is the set of leaves of Ψ. By Claim 3,

$$|S| + 2 \cdot \kappa(\Psi) = |V(\Psi)| \leq |T_1| + 2 \cdot \kappa(\Psi) - 2 \Rightarrow |S| \leq |T_1| - 2.$$

Therefore, (5.2) holds and this completes the proof of the lemma. □

5.2 Bounding the number and size of segment classes

In this section the goal is to prove the results of section 3.4 of [2] in our pd-graph setting. We introduce the notion of segment class that partitions the segments into classes of mutually "parallel'' segments. Using the results of the previous sections show that the number of these classes is again bounded by a linear function of the order of the pd-linkage.
As in the previous section, the special structure of a pd-graph is not of great importance for the proofs and when it is we will stretch it.

Classes of segments. Let G be a pd-graph and let $\mathcal{Q} = (\mathcal{C}, L)$ be a convex CL-configuration of G. Let S_1, S_2 be two segments of $\mathcal{Q}$ and let P and P' be two paths that are subgraphs of C_r, connect one endpoint of S_1 with an endpoint of S_2, and pass through no other endpoint of S_1 or S_2. We say that S_1 and S_2 are *parallel*, and we write $S_1 \parallel S_2$, if

(1) no segment of $\mathcal{Q}$ has both endpoints on P.

(2) no segment of $\mathcal{Q}$ has both endpoints on P'.

(3) the closed-interior of the cycle $P \cup S_1 \cup P' \cup S_2$ does not contain the disk D_0.

A *class of segment* is an equivalence class of segments of $\mathcal{Q}$ under the relation $\parallel$.

Given a pd-linkage L of G and a closed disk D of $\mathbb{R}^2$, we define $\mathbf{in}_D(L)$ to be the graph obtained from $L \cap D$ after dissolving all vertices that do not belong in L. Notice that $\mathbf{in}_{D_r}(L)$ is a biconnected outerplanar graph formed by the segments of $\mathcal{Q}$ and what remains from C_r after dissolving all vertices that do not belong in L. As $\mathcal{Q}$ is convex, one of the faces of $\mathbf{in}_{D_r}(L)$ contains the interior of D_0 and we call this face *central* face. We define the *segment tree* of $\mathcal{Q}$, denoted by $T(\mathcal{Q})$, as follows.

- let T^- be the weak dual of $\mathbf{in}_{D_r}(L)$ rooted at the vertex that corresponds to its central face.
- Let Q be the set of leaves of T^-. For each vertex $l \in Q$ do the following: Notice first that l is the dual of a face l^* of $\mathbf{in}_{D_r}(L)$. Let $W_1, \ldots, W_{\rho_l}$ be the extremal segments in the boundary of l^* (notice that, by the convexity of $\mathcal{Q}$, for every l, $\rho_l \geq 1$). Then, for each $i \in \{1, \ldots, \rho_l\}$, create a new leaf w_i corresponding to the extremal segment W_i and make it adjacent to l.

The *height* of $T(\mathcal{Q})$ is the maximum distance from its root to its leaves. The *real height* of $T(\mathcal{Q})$ is the maximum number of internal vertices of degree at least 3 in a path from its root to its leaves plus one. The *dilation* of $T(\mathcal{Q})$ is the maximum length of a path all whose internal vertices have degree 2 and are different from the root. For an example see Figure 5.2.

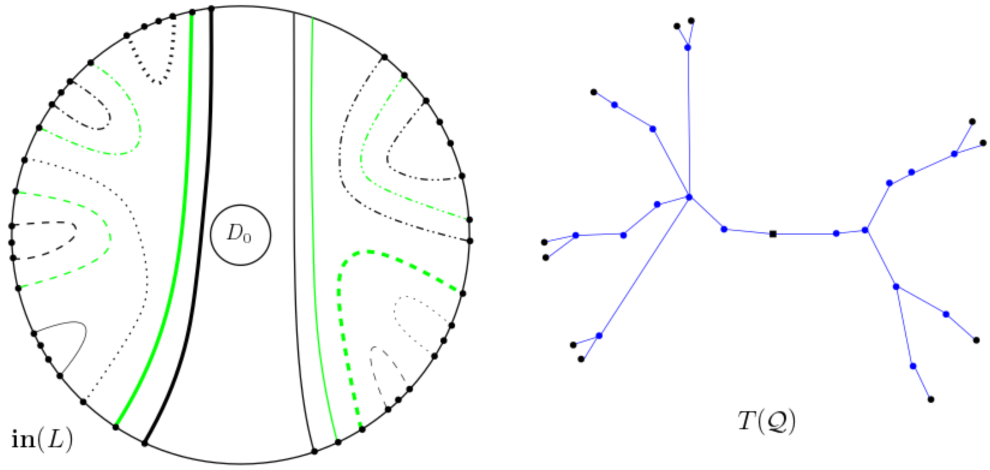

Figure 5.2: At the left the graph $\mathbf{in}_{D_r}(L)$ for some convex CL-configuration $\mathcal{Q} = (\mathcal{C}, L)$ and at the right the corresponding segment tree $T(\mathcal{Q})$. The CL-configuration $\mathcal{Q}$ has 11 extremal segments. There are 11 segment classes under the $||$ relation and internal edges of the same lining correspond to segments of the same class, while black (resp. green) internal edges correspond to primal (resp. dual) paths of L. At the right, the black square corresponds to the root of tree $(\mathcal{Q})$ and the black dots to its leaves. The dilation of $(\mathcal{Q})$ is 3, its height is 6 and its real height is 3.

Observation 5.2.1. *Let G be a pd-graph and let $\mathcal{Q} = (\mathcal{C}, L)$ be a convex CL-configuration of G. Then the height of $T(\mathcal{Q})$ is upper bounded by the dilation of $T(\mathcal{Q})$ multiplied by the real height of $T(\mathcal{Q})$.*

The next lemma is a consequence of Lemma 5.1.1 and the definition of a segment tree. We demand $L \cap C_r$ to be non-empty to ensure that L intersects D_r and the segment tree $T(\mathcal{Q})$ can be defined.

Lemma 5.2.1. *Let G be a pd-graph and $\mathcal{Q} = (\mathcal{C}, L)$ be a touch-free CL-configuration of G where $\mathcal{C}$ is tight in G, L is $\mathcal{C}$-cheap, and $L \cap C_r \neq \emptyset$. Then $\mathcal{Q}$ is convex and the real height of the segment tree $T(\mathcal{Q})$ is at most $2 \cdot |\mathcal{P}(L)| - 3$.*

Proof. The convexity of $\mathcal{Q}$ follows directly from Lemma 4.2.2. We examine the non-trivial case where $T(\mathcal{Q})$ contains at least one edge. We first claim that $|\mathcal{P}(L)| \geq 2$. Assume to the contrary that L consists of a single path P. As $\mathcal{Q}$ is convex and $L \cap C_r \neq \emptyset$, $\mathcal{Q}$ has at least one extremal segment. Suppose now that $\mathcal{Q}$ has more than one extremal segments, all of which are connected components of $C_r \cap P$.

Let P_1 and P_2 be the closures of the connected components of $L \setminus D_r$ that contain the terminals of P. Let $p_i \in V(C_r)$ be the endpoint of P_i that is not a terminal, $i \in \{1, 2\}$. Let also P' be any path in C_r between p_1 and p_2. Notice now that $P_1 \cup P' \cup P_2$ a pd-linkage that is equivalent to L and cheaper, a contradiction to the fact that L is $\mathcal{C}$-cheap. Therefore we conclude that $\mathcal{Q}$ has exactly one extremal segment, which contradicts the fact that $\mathcal{Q}$ is touch-free. Thus, $|\mathcal{P}(L)| \geq 2$ as claimed.

Recall that by the construction of $T(\mathcal{Q})$ there is a 1--1 correspondence between the leaves of $T(\mathcal{Q})$ and the extremal segments of $\mathcal{Q}$. From Lemma 5.1.1, $T(\mathcal{Q})$ has at most $2 \cdot |\mathcal{P}(L)| - 2$ leaves. Also $T(\mathcal{Q})$ has at least 2 leaves, because $\mathcal{Q}$ is touch-free. It is known that the number of internal vertices of degree ≥ 3 in a tree with $r \geq 2$ leaves is at most $r - 2$. Therefore, $T(\mathcal{Q})$ has at most $2 \cdot |\mathcal{P}(L)| - 4$ internal vertices of degree ≥ 3. Therefore the real height of $T(\mathcal{Q})$ is at most $2 \cdot |\mathcal{P}(L)| - 3$. □

5.3 Tidy pd-grids in convex configurations

In this section we prove that if a class of a segment in a CL-configuration of a pd-graph is "big" then there exists a "big" grid-like structure whose paths alternate between primal and dual. This is the counterpart of section 3.5 of [2] but here the primal-dual structure of a pd-graph and the types of the paths of a pd-linkage play a crucial role.

pd-patterns. A triple (H, χ, T) where H is a plane graph, $T \subseteq V(H)$, $\chi : V(H) \rightarrow \{p, d, c\}$ is called a *pd-pattern* if for any edge $e = \{x, y\} \in E(H)$, either $\{\chi(x), \chi(y)\} = \{p, c\}$ or $\{\chi(x), \chi(y)\} = \{d, c\}$. If $T = \emptyset$ we will just write (H, χ) to refer to the pd-pattern (H, χ, T). We will treat pd-patterns as structures embedded in $\mathbb{S}^2$.
Given an edge $e = \{x, y\} \in E(H)$, we say that it is a *primal edge* (rep. *dual edge*) of (H, χ, T) if $\chi(\{x, y\}) \setminus \{c\} = \{p\}$ (resp. $\chi(\{x, y\}) \setminus \{c\} = \{d\}$) and given a path P of G, we say that P is a *primal* (resp. *dual*) *path* of (H, χ, T) if all the edges of P are primal (resp. dual) edges.
Let G be a pd-graph. We say that G corresponds to the pd-pattern $(G, \psi_G, \emptyset)$, where $\psi_G : V(G) \rightarrow \{p, d, c\}$ and for every $\zeta \in \{p, d, c\}$ and every $v \in V(G)$:

$$\psi_G(v) = \zeta \quad \text{if and only if} \quad v \in V_\zeta(G)$$

Let $X \subseteq V(G)$. We call the pair (G, X) a *rooted pd-graph* and we say that X is its *root*. The rooted pd-graph (G, X), naturally corresponds to the pd-pattern (G, ψ_G, X).

pd-topological minors. Let (G, ψ, T) and (H, χ, Y) be two pd-patterns and $\lambda : Y \rightarrow T$ be a bijection. We say that the pd-pattern (H, χ, Y) is a λ-*pd-topological minor* of the pd-pattern (G, ψ, T) if there exists an injective function $\phi_0 : V(H) \rightarrow V(G)$ and a function $\phi_1 : E(H) \rightarrow \mathcal{P}(G)$ such that

- for every $x \in V(H)$, $\psi(\phi_0(x)) = \chi(x)$.
- $\lambda \subseteq \phi_0$.
- for every edge $e = \{x, y\} \in E(H)$, $\phi_1(\{x, y\})$ is a primal (resp. dual) path between $\phi_0(x)$ and $\phi_0(y)$ in (G, χ, T) if and only if e is a primal (resp. dual) edge of (H, ψ, Y).

- if two paths in $\phi_1(E(H))$ have a common vertex, then this vertex should be an endpoint of both paths.

Given the pair (ϕ_0, ϕ_1), we say that (H, χ, Y) is a *λ-pd-topological minor of (G, χ, T) via (ϕ_0, ϕ_1)* and we write $(H, \chi, Y) \leq^{\lambda}_{tm} (G, \psi, T)$. If $Y = T = \emptyset$ we will just say that (H, χ) is a *pd-topological minor* of (G, χ) and will write $(H, \chi) \leq_{tm} (G, \psi)$.
Let G is a pd-graph and $X \subseteq V(G)$. If a pd-pattern (H, χ, Y) is a λ-pd-topological minor of the pd-pattern (G, ψ_G, X) (the pd-pattern that corresponds to the rooted pd-graph (G, X)) via (ϕ_0, ϕ_1), we will also say that the pd-pattern (H, χ, Y) is a λ-pd-topological minor of the rooted pd-graph (G, X) via (ϕ_0, ϕ_1) and we will write $(H, \chi, Y) \leq^{\lambda}_{tm} (G, X)$.

pd-grids. Let G be a plane graph, k, k' be two integers, and Γ be a $(k \times k')$-grid of G. Let also $\chi : V(\Gamma) \to \{p, d, c\}$ be a 3-coloring of the vertices of Γ. We say that (Γ, χ) is a $(k \times k')$-*pd-grid* if the followings hold:

- In any horizontal and vertical line of Γ, the colors of its vertices alternate either between p and c (we call such a line a *primal line*) or between d and c (we call such a line a *dual line*).
- If the i-th horizontal (or horizontal) line of Γ is a primal (resp. dual) line, then its $(i + 1)$-th horizontal (resp. vertical) line is a dual (resp. primal) line.

A path P in (Γ, χ) is called a *primal* (resp. *dual*) *path* of (Γ, χ), if $V(P) \cap \chi^{-1}(d) = \emptyset$ (resp. $V(P) \cap \chi^{-1}(p) = \emptyset$). A cycle in (Γ, χ) is called a *primal* (resp. *dual*) *cycle* of (Γ, χ) if it is the union of two primal (resp. dual) paths of (Γ, χ)
For every pair of vertices (v_1, v_2) in $V(\Gamma)$, we say that (v_1, v_2) is a *primal* (resp. *dual*) *pair* if $(\chi(v_1), \chi(v_2)) \in \{(p, c), (c, p)\}$ (resp. $(\chi(v_1), \chi(v_2)) \in \{(d, c), (c, d)\}$. If $(v_1, v_2) = (c, c)$, then this pair can be considered to be either primal or dual. If (v, u) is a primal (resp. dual) pair and $\{v, u\} \in E(\Gamma)$, then we say that $\{v, u\}$ is a *primal* (resp. *dual*) *edge* of Γ (for an example of a pd-grid see Figure 5.3).
Let H be a pd-graph, Γ be a $(k \times k')$ grid of H, and let $\chi_H : V(H) \to \{p, d, c\}$ be such that for any $\pi \in \{p, d, c\},\ \chi_H(v) = \pi$ if and only if $v \in V_\pi(H)$. We say that Γ is a *pd-grid of H* if and only if (Γ, χ_H) is a pd-grid. Sometimes we will just use Γ to refer to the pd-grid (Γ, χ_H) if it is clear that χ_H is related to the types of vertices in $V(H)$.

Tilted pd-grids and L-tidy tilted pd-grids. Let G be pd-graph. A *tilted pd-grid* of G is a pair $\mathcal{U} = (\mathcal{X}, \mathcal{Z})$ where $\mathcal{X} = \{X_1, \ldots, X_r\}$ and $\mathcal{Z} = \{Z_1, \ldots, Z_r\}$ are collections of r primal and dual paths of G that are vertex-disjoint and such that

- X_1, X_r and Z_1, Z_r are primal paths of G and the types of $X_2, \ldots, X_r - 1$ and $Z_2, \ldots, Z_r - 1$ alternate between dual and primal.
- For each $i, j \in \{1, \ldots, r\}$ $I_{i,j} = X_i \cap Z_j$ is a (possibly edgeless) path of G and $I_{i,j}$ can contain an edge only if X_i and Z_j are of the same type.

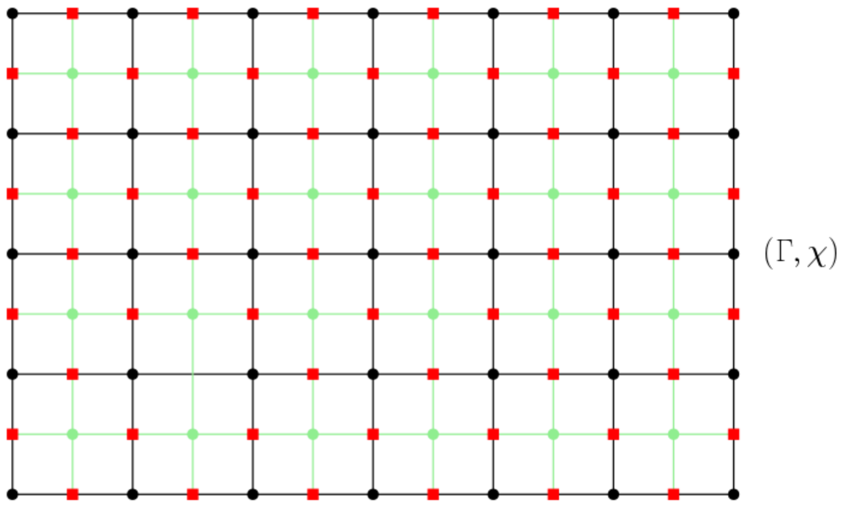

Figure 5.3: A pd-grid (Γ, χ) is depicted. The black dots correspond to the set $\chi^{-1}(p)$, the green dots to the set $\chi^{-1}(d)$, and the red squares to the set $\chi^{-1}(c)$. Any path (or cycle) that contains only black dots (resp. green dots) and red squares is a primal (resp. dual) path (or cycle) of the depicted pd-grid. Any edge with endpoints a black dot (resp. green dots) and a red square is a primal (resp. dual) edge of (Γ, χ).

- for $i \in \{1, \ldots, r\}$ the subpaths $I_{i,1}, I_{i,2}, \ldots, I_{i,r}$ appear in this order in X_i.
- for $j \in \{1, \ldots, r\}$ the subpaths $I_{1,j}, I_{2,j}, \ldots, I_{r,j}$ appear in this order in Z_j.
- $E(I_{1,1}) = E(I_{1,r}) = E(I_{r,r}) = E(I_{r,1}) = \emptyset$,
- Let
$$G_{\mathcal{U}} = (\bigcup_{i \in \{1,\ldots,r\}} X_i) \cup (\bigcup_{i \in \{1,\ldots,r\}} Z_i)$$
and let $G^*_{\mathcal{U}}$ be the graph taken from $G_{\mathcal{U}}$ after applying primal-dual contraction to all pairs of vertices that are adjoined through some vertex in $V(G^*_{\mathcal{U}}) \cap V_c(G)$ (all the crossing vertices on the paths). Then $G^*_{\mathcal{U}}$ contains an $(r \times r)$-pd-grid (Γ, χ) as a pd-topological minor.

We call the subgraph $G_{\mathcal{U}}$ of G *realization* of the tilted pd-grid $\mathcal{U}$ and the graph $G^*_{\mathcal{U}}$ *representation* of $\mathcal{U}$. We refer to the cardinality r of $\mathcal{X}$ (or $\mathcal{Z}$) as the *capacity* of $\mathcal{U}$. The perimeter of $G_{\mathcal{U}}$ is the (primal) cycle $X_1 \cup Z_1 \cup X_r \cup Z_r$. Given a pd-graph G and a pd-linkage L of G, we say that a tilted pd-grid $\mathcal{U} = (\mathcal{X}, \mathcal{Z})$ of G is an *L-tidy tilted pd-grid* of G if $\mathcal{D}_{\mathcal{U}} \cap L = \bigcup\mathcal{Z}$, where $\mathcal{D}_{\mathcal{U}}$ is the closed interior of the perimeter of $G_{\mathcal{U}}$.

Lemma 5.3.1. *Let G be a pd-graph and let $\mathcal{Q} = (\mathcal{C}, L)$ be a convex CL-configuration of G. Let also $\mathcal{S}$ be an equivalence class of the relation $\|$. Then G contains a tilted pd-grid $\mathcal{U} = (\mathcal{X}, \mathcal{Z})$ of capacity $\lceil \frac{|\mathcal{S}|-2}{2} \rceil$ that is an L-tidy tilted pd-grid of G.*

Proof. Let $\mathcal{C} = \{C_0, \ldots, C_r\}$ and let $\mathcal{S} = \{S_1, \ldots, S_m\}$. We assume that S_1 and S_m are primal paths of G. For each $i \in \{1, \ldots, m\}$, let σ_i be the eccentricity of S_i and let $\sigma^{\max} = \max\{\sigma_i \mid i \in \{1, \ldots, m\}\}$ and $\sigma^{\min} = \min\{\sigma_i \mid i \in \{1, \ldots, m\}\}$. Convexity allows us to assume that $S_1, \ldots, S_m$ are ordered in a way that

- The types of the paths in $\mathcal{S}$ alternate between primal and dual.
- $\sigma_1 = \sigma^{\min}$,
- $\sigma_m = \sigma^{\max}$, and
- for all $i \in \{1, \ldots, m-1\}$, $\sigma_{i+1} = \sigma_i + 1$.
- for all $i \in \{1, \ldots, m\}$, $I_{i,\sigma_i} = S_i \cap C_{\sigma_i}$ is a (possibly edgeless) subpath of C_{σ_i}.

Let $m' = \lceil \frac{m}{2} \rceil$ and let x, x' (resp. y, y') be the endpoints of the path S_1 (resp. $S_{m'}$) such that the one of the two (x, y)-paths (resp. (x', y')-paths) in C_r contains both x', y' (x, y) and the other, say P (resp. P'), contains none of them. Let D_S be the closed-interior of the cycle $S_1 \cup P' \cup S_{m'} \cup P$. Let also A_C be the union of the intersection of G with $D_{\sigma^{\max}-(m'-1)} \setminus D_{\sigma^{\max}}$ and the cycles $C_{\sigma^{\max}-(m'-1)}$ and $C_{\sigma^{\max}}$. Let Δ be any of the two connected components of $D_S \cap A_C$. We now consider the graph

$$(L \cup \mathbf{U}\mathcal{C}) \cap \Delta.$$

It is now easy to verify that the above graph is the realization $G_{\mathcal{U}}$ of a tilted pd-grid $\mathcal{U} = (\mathcal{X}, \mathcal{Z})$ of capacity m', where the paths in $\mathcal{X}$ are the portions of the cycles $C_{\sigma^{\max}-(m'-1)}, \ldots, C_{\sigma^{\max}}$ cropped by Δ, while the paths in $\mathcal{Z}$ are the portions of the paths in $\{S_1, \ldots, S_{m'}\}$ cropped by Δ. As $\mathcal{S}$ is an equivalence class of $\|$, it follows that $\mathcal{U}$ is L-tidy.
At the beginning of the proof we assumed that S_1 and S_m are primal paths. If this is not the case we can discard them and repeat the same argument for $\mathcal{S}' \setminus \{S_1, S_m\}$. In any case, we obtain a tilted pd-grid of the required capacity. □

CHAPTER 6

INSIDE THE OUTER CYCLE

6.1 Replacing pd-linkages by cheaper ones

In this section we prove that a pd-linkage of order k can be rerouted to a cheaper one, given the existence of an L-tidy tilted pd-grid of capacity greater at least $2^k + 1$. In other words, the existence of a cheap pd-linkage L implies an exponential (on the order of L) upper bound on the capacity of any L-tidy tilted pd-grid.

Let G be a pd-graph and let L be a pd-linkage in G. Let also D be a closed disk in the surface where G is embedded. We say that L *crosses vertically* D if the outerplanar graph defined by the boundary of D and $L \cap D$ has exactly two simplicial faces. Hence, the vertices of **bnd**$(D) \cap L$ are naturally partitioned into the *up* and *down* ones.

We will need the following two main lemmas. The first one is exactly Lemma 2 from [1] (the "reflection trick") and the proof of the second one is a variance of the proof of a claim stated in page 11 of [1], slightly transformed to fit in our context. Lemma 6.1.3 then follows easily and is crucial for the remaining part.

Lemma 6.1.1. *Let Σ be an alphabet of size $|\Sigma| = k$, and let $w \in \Sigma^*$ be a word over Σ of length $|w| > 2^k$. Then, w contains an infix y such that every letter occurring in y, occurs an even number of times.*

Proof. Let $\Sigma = \{a_1, \ldots, a_k\}$ and let $w = w_1 \ldots w_n$ where $n > 2k$. We define the vectors $z_i \in \{0, 1\}^k$ for every $i \in \{1, \ldots, n\}$, and we let the jth entry of z_i be 0 if and only if a_j occurs an even number of times in the prefix $w_1 \ldots w_i$ of w and 1 otherwise. Since $n > 2k$, there exist two indices $i, i' \in \{1, \ldots, n\}$ with $i \neq i'$, such that $z_i = z_{i'}$. Let $a_l \in \Sigma$ that appears in the infix $w_{i+1}w_{i+2} \ldots w_{i'}$. Then, as the lth coordinate of z_i is the same as the lth coordinate of $z_{i'}$, a_l appears in $w_{i+1}w_{i+2} \ldots w_{i'}$ an even number of times. Thus, the infix $w_{i+1}w_{i+2} \ldots w_{i'}$ satisfies the requirements of the statement. $\square$

Lemma 6.1.2. *Let G be a plane graph and let D be a closed disk and a linkage L of G of order k that crosses D vertically. Let also $L \cap D$ consist of $r > 2^k$ lines. Then, there is a collection $\mathcal{N}$ of strictly less than r mutually non-crossing lines in D each connecting two points of* **bnd**$(D) \cap L$, *such that there exists some linkage R that is a subgraph of $L \setminus$* **int**(D) *such that $R \cup \mathbf{U}\mathcal{N}$ is a linkage of the graph $(G \setminus D) \cup \mathbf{U}\mathcal{N}$ that is equivalent to L.*

Proof. Let $\mathcal{P}(L) = \{P_1, \ldots, P_k\}$. Every line in $L \cap D$ is a subpath of exactly one path in $\mathcal{P}(L)$. Let $\Sigma = \{a_1, a_2, \ldots, a_k\}$ be an alphabet of size k, where a_i corresponds to the path P_i.
Then, as the lines cross D vertically, there is an ordering that indicates the way that they consecutively appear in D. We can naturally map such an ordering to a word, say w, over Σ by replacing every line l_j in the ordering by a_i if l_i is a subpath of P_i.
Observe that lemma 6.1.1 can be applied for Σ and w, therefore we obtain that there is an infix y of w such that every letter occurring in it, occurs an even number of times. This, "translated" back to lines, implies that there is a non-empty subset $A \subseteq L \cap D$ of lines that appear consecutively in D and for every $P_i \in \mathcal{P}(L)$, the number of lines in A that are subpaths of P_i is even (it can be zero).
Let $A = \{l_1, \ldots, l_{|y|}\}$. For every path $P_i \in \mathcal{P}(L)$, we define $A_{P_i} = \{l \in A \mid l \text{ is a subpath of } P_i\}$ and we know that $|A_{P_i}|$ is even. Therefore, for some $n_i \in \mathbb{N}$, $A_{P_i} = \{l_1^i, \ldots, l_{2n_i}^i\} \subseteq A$, for every $P_i \in \mathcal{P}(L)$ for which $A_{P_i} \neq \emptyset$.
For every such i, we traverse P_i from s_i to t_i and orient the lines of A in the way that we meet them. For every odd number $j \in \{1, \ldots, 2n_i\}$ we replace the subpath of P_i from $tail(l_j^i)$ to $head(l_{j+1}^i)$ by a new line f_j^i which lies in D avoiding crossings (for an example see Figure 6.1).
After having done these replacements we obtain a new path, P_i', from s_i to t_i that contains strictly less lines than P_i in A, and as this operation only causes changes in A, also in D.
Let B be the set containing the lines introduced from the replacement. Then, if we set $\mathcal{N} = ((L \cap D) \cup B) \setminus A$ it is easy to observe that there exists a linkage R that is a subgraph of $L \setminus$ **int**(D), such that $R \cup \mathbf{U}\mathcal{N}$ meets the requirements of the statement. □

Lemma 6.1.3. *Let (G, χ) be a pd-pattern, D be a closed disk, and L a pd-linkage of order k in G that crosses D vertically. Let also $L \cap D$ consist of $r > 2^k$ lines. Then, there is a collection $\mathcal{N}$ of strictly less than r mutually non-crossing lines in D each connecting two points of* **bnd**$(D) \cap L$, *such that there exists some pd-linkage R that is a subgraph of $L \setminus$* **int**(D) *such that $R \cup \mathbf{U}\mathcal{N}$ is a pd-linkage of the graph $(G \setminus D) \cup \mathbf{U}\mathcal{N}$ that is equivalent to L.*

Proof. As L is a pd-linkage of the pd-graph G, every path $P \in \mathcal{P}(L)$ is either a primal or a dual path of G.
Let R be the linkage provided by Lemma 6.1.2. Then, as R is a subgraph of $L \setminus$ **int**(D), every path of L that is also a path of R, preserves its type.
As for the lines in B that are introduced in D after the replacement, they inherit their types from the lines that they replace. □

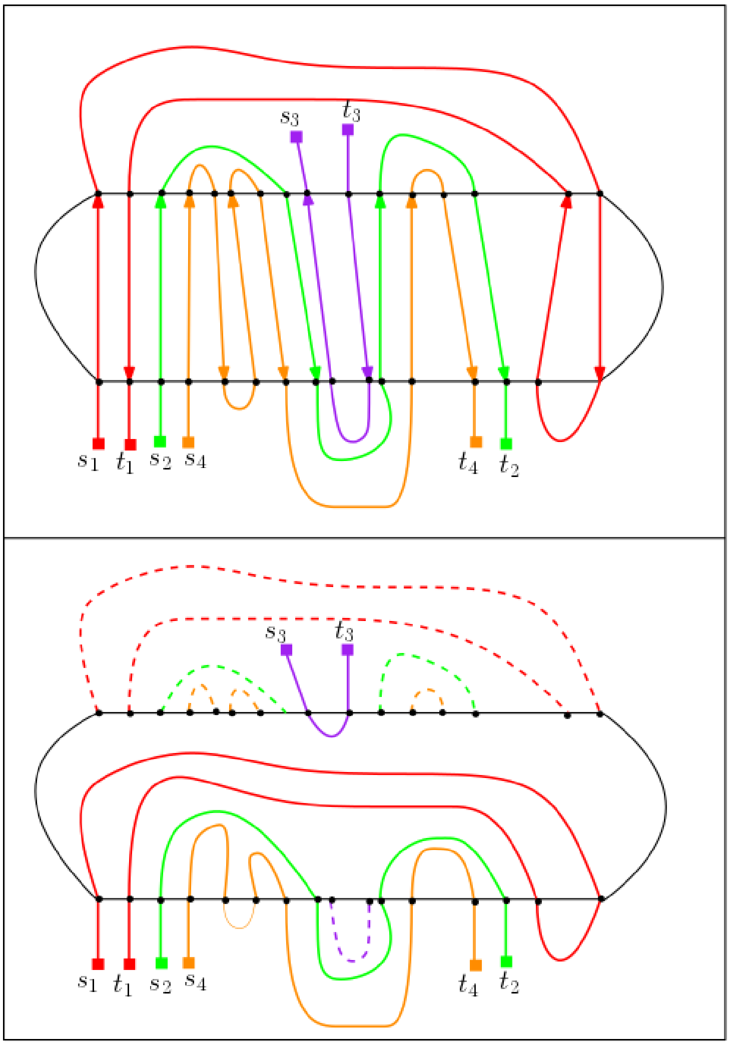

Figure 6.1: A visualization of the proof of the "reflection trick" in Lemma 6.1.2

Lemma 6.1.4. *Let k, k', ρ be integers such $\rho \leq \min\{k, k'\}$. Let (Γ, χ) be a pd-grid whose perimeter is a primal cycle, where Γ is a $(k \times k')$-grid, and let $\{v_1^{\text{up}}, \ldots, v_\rho^{\text{up}}\}$ (resp. $\{v_1^{\text{down}}, \ldots, v_\rho^{\text{down}}\}$) be vertices of the higher (resp. lower) horizontal line arranged as they appear in it from left to right and such that $(v_h^{\text{up}}, v_h^{\text{down}})$ is either a primal or a dual pair of (Γ, χ) (we consider all pairs with both vertices in $\chi^{-1}(c)$ to be dual), for every $h \in [\rho]$. Then (Γ, χ) contains ρ pairwise disjoint paths $P_1, \ldots, P_\rho$ such that, for every $h \in [\rho]$, P_h is a primal (resp. dual) path with endpoints v_h^{up} and v_h^{down} iff $(v_h^{\text{up}}, v_h^{\text{down}})$ is a primal (resp. dual) pair of (Γ, χ).*

Proof. We use induction on ρ. The statement trivially holds for $\rho = 0$. Let $(i, j) \in [k]^2$ such that v_ρ^{up} (resp. v_ρ^{down}) is the i-th (resp. j-th) vertex of the higher (lower) horizontal line counting from left to right. We first consider the

case where $i \geq j$. Let P_ρ be the path created by starting from v_ρ^{up}, moving $k'-1$ edges down (observe that the vertices we meet have the required color), and then $j-i$ edges to the left. For $h \in [\rho-1]$ let $P_h^{(\text{down})'}$ be the path created by starting from v_h^{down} and moving one edge up (clearly, $P_h^{(\text{down})'}$ consists of a single edge). We also denote by $v_h^{(\text{down})'}$ the other endpoint of $P_h^{(\text{down})'}$. We now define Γ' as the subgrid of Γ that occurs from Γ after removing its lower horizontal line and, for every $h \in [i,k]$, its h-th vertical line. It is clear that (Γ', χ') is a pd-grid, where χ' is the restriction of χ to the vertices of Γ'. Observe also that for every $h \in [\rho-1]$, the pair $(v_h^{\text{up}}, v_h^{\text{down}})$ is substituted by the pair $(v_h^{\text{up}}, v_h^{(down)'})$ which we consider to be primal (resp. dual) iff $(v_h^{\text{up}}, v_h^{\text{down}})$ is primal (in order to maintain the initial requirements about the types of the paths). By construction, none of the edges or vertices of P_ρ belongs to Γ'. Notice also that the higher (resp. lower) horizontal line of Γ' contains all vertices in $\{p_1^{\text{up}}, \dots, p_{\rho-1}^{\text{up}}\}$ (resp. $\{p_1^{(\text{down})'}, \dots, p_{\rho-1}^{(\text{down})'}\}$) and by the induction hypothesis, (Γ', χ') contains the $\rho-1$ pairwise disjoint paths $P_1', \dots, P_{\rho-1}'$ that meet the conditions of the statement.
It is now easy to verify that $P_1' \cup P_1^{(\text{down})'}, \dots, P_{\rho-1}' \cup P_{\rho-1}^{(\text{down})'}, P_\rho$ is the required collection of pairwise disjoint paths.
For the case where $i < j$, we can just think of grid Γ being turned upside down and repeat the same argument. □

Lemma 6.1.5. *Let k be an odd integer and (Γ, χ) be a pd-grid, where Γ is a $(k \times 2k)$-grid embedded in the plane and assume that the vertices of its outer cycle, arranged in clockwise order, are*

$$\{v_1^{\text{up}}, \dots, v_k^{\text{up}}, v_2^{\text{right}}, \dots, v_{k-1}^{\text{right}}, v_k^{\text{down}}, \dots, v_1^{\text{down}}, v_{k-1}^{\text{left}}, \dots, v_2^{\text{left}}, v_1^{\text{up}}\}$$

Additionally, $\chi(v_1^{\text{up}}) = p$ and for every $v \in V(\Gamma) \setminus \{v_1^{\text{up}}\}$ such that $v' \in V(\Gamma)$ precedes v in the cyclic arrangement of $V(\Gamma)$, $\chi(v) = c$ if $\chi(v') = p$ and $\chi(v) = p$ otherwise.
Let also (H, ψ) be a pd-pattern. The vertices of H have degree 0 or 1 and can be cyclically arranged in clockwise order as

$$\{x_1^{\text{up}}, \dots, x_k^{\text{up}}, x_k^{\text{down}}, \dots, x_1^{\text{down}}, x_1^{\text{up}}\}$$

such that if we add to H the edges formed by pairs of consecutive vertices in this cyclic ordering, the resulting graph H^+ is outerplanar and $\psi(x_i^{\text{up}}) = \chi(v_i^{\text{up}})$, $\psi(x_i^{\text{down}}) = \chi(v_i^{\text{down}})$ for every $i \in \{1, \dots, k\}$. Also, for every edge $e = \{v, u\} \in E(H)$, either $\psi(v) = \psi(u) = p$ (primal edges of H) or $\psi(v) = \psi(u) = c$ (dual edges of H). Let V^1 be the vertices of H that have degree 1 and let $H^1 = H[V^1]$. Then $(H^1, \psi|_{V^1})$ is a pd-topological minor of (Γ, χ) via some pair (ϕ_0, ϕ_1), satisfying the following properties:

1. *$\phi_0(x_i^{\text{up}}) = v_i^{\text{up}}$, $i \in \{1, \dots, k\} \cap V^1$*
2. *$\phi_0(x_i^{\text{down}}) = v_i^{\text{down}}$, $i \in \{1, \dots, k\} \cap V^1$.*

Proof. Let $U = \{x_1^{\text{up}}, \dots, x_k^{\text{up}}\} \cap V^1$ and $D = \{x_1^{\text{down}}, \dots, x_k^{\text{down}}\} \cap V^1$. Let ϕ_0 be as required. In the rest of the proof we provide the definition of ϕ_1. We partition the edges of H^1 into three sets: the *upper edges* E_U that connect vertices in U, the *down edges* E_D that connect vertices in D, and the *crosssing edges* E_C that have one endpoint in U and one in D. As $|V(H^1)| \leq 2k$ we have that $|E(H^1)| \leq k$ and therefore $|E_U|+|E_D|+|E_C| = |E(H^1)| \leq k$. We set $\rho = |E_C|$.

We recursively define the *depth* of an edge $e = \{x_i^{\text{up}}, x_j^{\text{up}}\}$ in E_U as follows:

- If e is a dual (resp. primal) edge of H, then its depth is 0 (resp. 1) if there is no edge of E_U with an endpoint in $\{x_{i+1}^{\text{up}}, \ldots, x_{j-1}^{\text{up}}\}$.
- If the maximum depth of an edge with an endpoint in $\{x_{i+1}^{\text{up}}, \ldots, x_{j-1}^{\text{up}}\}$ is i and this edge is primal, respectively dual, then the depth of e is $i+2$, respectively $i+1$.

The depth of an edge $e' = \{x_i^{\text{down}}, x_j^{\text{down}}\}$ is defined analogously. It follows, by the definition (observe that a worst case scenario is realized when E_U (resp. E_D) is the set of all primal vertices in $\{x_1^{\text{up}}, \ldots, x_k^{\text{up}}\}$ (resp. $\{x_1^{\text{down}}, \ldots, x_k^{\text{down}}\}$)) that:

$$q^{\text{up}} = \max\{\textit{depth}(e) \mid e \in E_U\} + 1 \quad \leq \quad 2|E_U| \tag{6.1}$$

$$q^{\text{down}} = \max\{\textit{depth}(e) \mid e \in E_D\} + 1 \quad \leq \quad 2|E_D| \tag{6.2}$$

We now define $\phi_1 : E(H) \to \mathcal{P}(\Gamma)$ as follows:

- For every edge $e = \{x_i^{\text{up}}, x_j^{\text{up}}\}$ in E_U of depth l and such that $i < j$
  - If $\psi(x_i^{\text{up}}) = \psi(x_j^{\text{up}}) = p$, i.e. e is a primal edge, let $\phi_1(e)$ be the path defined if we start in the grid Γ from v_i^{up}, move $2l$ steps down, then $j-i$ steps to the right, and finally $2l$ steps up to the vertex v_j^{up} (the term "number of steps" refers to the number of edges traversed). Observe that the obtained path is a primal path of Γ.
  - if $\psi(x_i^{\text{up}}) = \psi(x_j^{\text{up}}) = c$, i.e. e is a dual edge, let $\phi_1(e)$ be the path defined if we start in the grid Γ from v_i^{up}, move $2l-1$ steps down, then $j-i$ steps to the right, and finally $2l-1$ steps up to the vertex v_j^{up}. Observe that in this case, the obtained path is a dual path of Γ.
- For every edge $e = \{x_i^{\text{down}}, x_j^{\text{down}}\}$ in E_D of depth l and such that $i < j$
  - If $\psi(x_i^{\text{up}}) = \psi(x_j^{\text{up}}) = p$, i.e. e is a primal edge, let $\phi_1(e)$ be the path defined if we start in the grid Γ from v_i^{up}, move $2l$ steps up, then $j-i$ steps to the right, and finally $2l$ steps down to the vertex v_j^{up}. Observe that the obtained path is again a primal path of Γ.
  - if $\psi(x_i^{\text{down}}) = \psi(x_j^{\text{down}}) = c$, i.e. e is a dual edge, let $\phi_1(e)$ be the path defined if we start in the grid Γ from v_i^{up}, move $2l-1$ steps up, then $j-i$ steps to the right, and finally $2l-1$ steps down to the vertex v_j^{up}. Observe that the obtained path is again a dual path of Γ.

We have defined the value of ϕ_1 for all edges in $E_U \cup E_D$ and it is easy to confirm that all paths in $\phi_1(E_U \cup E_D)$ are mutually non-crossing. It remains to define $\phi_1(e)$ for every $e \in E_C$.

Notice that the distance between $\phi_0(U)$ and some horizontal line of Γ that contains edges of the images of the upper edges is $2 \cdot \max\{\textit{depth}(e) \mid e \in E_U\}$ that, from (6.1), is equal to $2q^{\text{up}} - 2$. Symmetrically, using (6.2) instead of (6.1), the distance between $\phi_0(D)$ and the horizontal lines of Γ that contain edges of the images of the down edges is equal to $2q^{\text{down}} - 2$. As a consequence, the graph

$$\hat{\Gamma} = \Gamma \setminus \{x \in V(\Gamma) \mid \textbf{dist}_\Gamma(x, \phi_0(U)) < 2q^{\text{up}} \ \vee \ \textbf{dist}_\Gamma(x, \phi_0(D)) < 2q^{\text{down}}\}$$

is a $(k \times k')$-grid, where $k' \geq 2k - 2(q^{\text{up}} + q^{\text{down}})$, whose vertices do not appear in any of the paths in $\phi_1(E_U \cup E_D)$. Given a crossing edge $e = \{x_i^{\text{up}}, x_j^{\text{down}}\} \in E_C$, we define the path P_e^{up} as the subpath of Γ created if we start from x_i^{up} and then go $2q^{\text{up}}$ steps down. Similarly, we define P_e^{down} as the subpath of Γ created if we start from x_j^{down} and then go $2q^{\text{down}}$ steps up. Notice that each of the paths P_e^{up} (resp. P_e^{down}) share only one vertex, say p_e^{up} (resp. p_e^{down}), with Γ' that is one of their endpoints. We use the notation $\{p_1^{\text{up}}, \ldots, p_\rho^{\text{up}}\}$ (resp. $\{p_1^{\text{down}}, \ldots, p_\rho^{\text{down}}\}$) for the vertices of the set $\{p_e^{\text{up}} \mid e \in E_C\}$ (resp. $\{p_e^{\text{down}} \mid e \in E_C\}$) such that, for every $h \in [\rho]$, there exists an $e \in E_C$ such that p_h^{up} is an endpoint of P_e^{up} and p_h^{down} is an endpoint of P_e^{down}. We also agree that the vertices in $\{p_1^{\text{up}}, \ldots, p_\rho^{\text{up}}\}$ (resp. $\{p_1^{\text{down}}, \ldots, p_\rho^{\text{down}}\}$) are ordered as they appear from left to right in the upper (lower) horizontal line of $\hat{\Gamma}$ (this is possible because of the outeplanarity of H^+).

Notice that $\rho = |E(H^1)| - (|E_U| + |E_D|) \leq k - (|E_U| + |E_D|)$ which by (6.1) and (6.2) implies that $\rho \leq k'$.

As $\rho \leq k' \leq k$, we can now apply Lemma 6.1.4 on $(\hat{\Gamma}, \chi|_{V(\hat{\Gamma})})$, $\{p_1^{\text{up}}, \ldots, p_\rho^{\text{up}}\}$ and $\{p_1^{\text{down}}, \ldots, p_\rho^{\text{down}}\}$ and obtain a collection $\{P_e \mid e \in E_C\}$ of ρ pairwise disjoint paths in $(\hat{\Gamma}\chi|_{V(\hat{\Gamma})})$ between the vertices of $\{p_e^{\text{up}} \mid e \in E_C\}$ and the vertices of $\{p_e^{\text{down}} \mid e \in E_C\}$, which respect the type of each pair. It is now easy to verify that $\{P_e^{\text{up}} \cup P_e \cup P_e^{\text{down}} \mid e \in E_C\}$ is a collection of ρ vertex disjoint paths between U and D. We can now complete the definition of ϕ_1 for the crossing edges of H by setting, for each $e \in E_C$, $\phi(e) = P_e^{\text{up}} \cup P_e \cup P_e^{\text{down}}$. By the above construction it is clear that the pd-pattern $(H^1, \psi|_V^1)$ is a pd-topological minor of the pd-grid (Γ, χ) via the pair of functions (ϕ_1, ϕ_2). For a visualization of the idea of the proof, see an example in Figure 6.1. □

Lemma 6.1.6. *Let G be a pd-graph, L be a pd-linkage of order k in G, and $\mathcal{U} = (\mathcal{X}, \mathcal{Z})$ be an L-tidy tilted pd-grid of G with capacity $2m$. Let also Δ be the closure of the perimeter of $G_{\mathcal{U}}$. If $2m > 2^{2k}$, then G contains a pd-linkage L' such that*

1. *L and L' are equivalent,*
2. *$L' \setminus \Delta \subseteq L \setminus \Delta$, and*
3. *$|E(\mathbf{U}\mathcal{Z} \cap L')| < |E(\mathbf{U}\mathcal{Z} \cap L)|$.*

Proof. Let $\mathcal{X} = \{X_1, \ldots, X_{2m}\}$ and $\mathcal{Z} = \{Z_1, \ldots, Z_{2m}\}$. Let also $G_{\mathcal{U}}$ be the realization of $\mathcal{U}$ in G and G^* (resp. L^*) be the pd-graph (resp. pd-linkage) obtained from G (resp. L) if we apply primal-dual contraction to all pair of vertices that are adjoined through a vertex in $V_c(G)$ that belongs to some path of $\bigcup_{(i,j)\in\{1,\ldots,r\}^2} I_{i,j}$, where $I_{i,j} = X_i \cap Z_j$, $i,j \in \{1,\ldots,m\}$. We also define $\mathcal{X}^*$ and $\mathcal{Z}^*$ by applying the same operation to their paths. Notice that $\mathcal{U}^* = (\mathcal{X}^*, \mathcal{Z}^*)$ is an L^*-tidy tilted pd-grid of G^* with capacity $2m$ and that the lemma follows if we find a linkage L'^* such that the three conditions of the statement are true for Δ^*, L^*, L'^*, and $\mathcal{Z}^*$, where Δ^* is the closed-interior of the perimeter of $G^*_{\mathcal{U}}$.

Let $G^{*-} = (G^* \setminus \Delta^*) \cup \mathbf{U}\mathcal{Z}$ and apply lemma 6.1.2 on G^{*-}, Δ^*, and L^* (the lemma can be applied as $2m > 2^{2k} > 2^k$). Let $\mathcal{N}$ be a collection of strictly less than $2m$ mutually non-crossing lines in D each connecting two points of $\mathbf{bnd}(\Delta^*) \cap L^*$ and a linkage $R \subseteq L^* \setminus \mathbf{int}(\Delta^*)$ such that $L_0 = R \cup \mathbf{U}\mathcal{N}$ is a pd-linkage of the graph $(G^* \setminus \Delta^*) \cup \mathbf{U}\mathcal{N}$ that is equivalent to L^*. Let $H = (L_0 \cap \Delta^*) \cup (L^* \cap \mathbf{bnd}(\Delta^*))$. Notice that in H, the set $V(L_0 \cap \Delta^*)$ contains the vertices of H of degree 1 while the rest of the vertices of H have degree 0 and all edges of H have their endpoints in $V(L_0 \cap \Delta^*)$. Recall that the $(2m \times 2m)$-pd-grid is a pd-topological minor of $G^*_{\mathcal{U}}$ and clearly the $(m \times 2m)$-pd-grid is a pd-topological minor of $G^*_{\mathcal{U}}$.

We can now apply Lemma 6.1.5 for the $(m \times 2m)$-grid Γ and H. We obtain that $H^1 = L_0 \cap \Delta^*$ is a pd-topological minor of Γ via some pair (ϕ_0, ϕ_1). We now define the graph

$$L = \bigcup_{e\in E(H^1)} E(\phi_1(e)).$$

Notice that L is a subgraph of Γ. We also define the graph

$$Q = \bigcup_{e\in E(L)} \phi_1(e)$$

which, in turn, is a subgraph of $G^*_{\mathcal{U}}$. Observe that $L'^* = R \cup Q$ is a pd-linkage of G^* that is equivalent to L^*. This proves Condition 1. Condition 2 follows from the fact that $R \subseteq L^* \setminus \mathbf{int}(\Delta^*)$. Notice now that, as $|\mathcal{N}| < 2m$, $E(\mathbf{U}\mathcal{Z}^* \cap Q)$ is a proper subset of $E(\mathbf{U}\mathcal{Z}^*)$. By construction of L'^*, it holds that $E(\mathbf{U}\mathcal{Z} \cap L'^*) = E(\mathbf{U}\mathcal{Z} \cap Q)$. Moreover, as $\mathcal{U}^* = (\mathcal{X}^*, \mathcal{Z}^*)$ is an L^*-tidy pd-tilted grid of G^*, it follows that $E(\mathbf{U}\mathcal{Z}^*) = E(\mathbf{U}\mathcal{Z}^* \cap L^*)$. Therefore, Condition 3 follows. ☐

6.2 Existence of an irrelevant crossing

We now bring together all results from previous sections in order to prove the main theorem of this thesis, that roughly states that if in an instance for PD-DP problem the treewidth of the pd-graph is sufficiently large, then there exists a pair of vertices that are adjoined, say through a vertex v, such that by applying primal-dual contraction to them we obtain an equivalent instance. This pair is an analogue of an irrelevant vertex in our context, and we call vertex v (that "connects" the vertices of the adjoined pair) an *irrelevant crossing*.

Lemma 6.2.1. *Let G be a plane graph containing a $(k \times k)$-grid Γ as a minor, $T \subseteq V(G)$, and $J \in$ **duals**(G). If $k \geq (r+1) \cdot \lceil \sqrt{|T|+1} \rceil$, then the pd-graph*

$\mathbf{pd}(G, J)$ contains an alternating tight sequence of r concentric cycles $\mathcal{H} = \{H_1, \ldots, H_r\}$ such that none of the vertices in the closed interior of H_r is in T.

Proof. We will first prove that G (and thus also $\mathbf{pd}(G, J)$) contains a tight sequence of $\lceil \frac{r+1}{2} \rceil$ concentric cycles (primal cycles in $\mathbf{pd}(G, J)$). From the definition of tightness,the existence of a tight sequence of $\lceil \frac{r+1}{2} \rceil$ concentric cycles in G follows if we just prove the existence of a sequence of $\lceil \frac{r+1}{2} \rceil$ concentric cycles in G. It is clear that Γ contains $|T|+1$ vertex disjoint $(r+1)\times(r+1)$-grids and thus G (and also $\mathbf{pd}(G, J)$) contains a $(r+1) \times (r+1)$-grid Γ' as a minor and this is certified by a function $\phi : V(\Gamma) \to 2^{V(G')}$ where G' is a subgraph of G that does not contain vertices of T. Notice that $V(\Gamma')$ can be partitioned into sets $\{V_0, \ldots, V_{\lceil \frac{r+1}{2} \rceil}\}$, corresponding to $\lceil \frac{r+1}{2} \rceil$ concentric cycles of Γ' that are arranged from inside to outside, i.e., V_0 contains the centers of Γ'. In each $G[\phi(V_i)]$, pick a cycle C_i meeting the models of all vertices of V_i. Because G is a plane graph, it is easy to verify that $\{C_1, \ldots, C_{\lceil \frac{r+1}{2} \rceil}\}$ is a collection of concentric cycles in G.
Fix now two consecutive cycles, say C_i and C_{i+1} with $i \leq \lceil \frac{r+1}{2} \rceil$. Then, observe that, as Γ' is a grid minor in G, there exist edges (that corresponds to disjoint paths of G) with one endpoint in $V(C_i)$ and the other in $V(C_{i+1})$, that partition $D_i \setminus D_{i+1}$ to open sets, which correspond to faces, say $f_1, \ldots f_l$, in the subgraph of G obtained after deleting all vertices of G in $\mathbf{int}(D_i \setminus D_{i+1})$. The vertices in $V_d(\mathbf{pd}(G, J))$ that correspond to these faces, induce a dual cycle $C_{i,i+1}$, such that $D_i \subset D_{i,i+1} \subset D_{i+1}$. If $r \equiv 1 \mod 2$, then $\mathcal{H} = \{C_1, C_{1,2}, C_2, \ldots, C_{\lceil \frac{r+1}{2} \rceil -1, \lceil \frac{r+1}{2} \rceil}, C_{\lceil \frac{r+1}{2} \rceil}\}$ and if $r \equiv 0 \mod 2$ we get an extra cycle. As we have previously observed, the dual cycles in $\mathcal{H}$ can be replaced such that the obtained sequence $\mathcal{H}'$ is an alternating tight sequence of concentric cycles in G. □

The next lemma follows from [21]:

Lemma 6.2.2. *Let G be a plane graph and let $J \in \mathbf{duals}(G)$. Then, $\mathbf{tw}(\mathbf{pd}(G, J)) \leq 2 \cdot \mathbf{tw}(G)$.*

From the last two lemmas and proposition 2.3.1 follows

Lemma 6.2.3. *Let G be a pd-graph such that $\mathbf{tw}(G) \geq 9 \cdot (r+1) \cdot \lceil \sqrt{|T|+1} \rceil$ and let $T \subseteq V(G)$. Then G contains an alternating tight sequence of r concentric cycles $\mathcal{C} = \{C_1, \ldots, C_r\}$ such that none of the vertices in the closed interior of H_r is in T.*

Proof. As $\mathbf{tw}(G) \geq 9 \cdot (r+1) \cdot \lceil \sqrt{|T|+1} \rceil$, Lemma 6.2.2 gives that $\mathbf{tw}(G^p) \geq 4.5 \cdot (r+1) \cdot \lceil \sqrt{|T|+1} \rceil$ (remember that G^p is the primal part of the pd-graph G). Moreover, Lemma 2.3.1 implies that G^p contains a $\big((r+1)\lceil \sqrt{|T|+1} \rceil\big) \times \big((r+1)\lceil \sqrt{|T|+1} \rceil\big)$-grid as a minor and as a result, by Lemma 6.2.1, we get that G contains an alternating tight sequence of r concentric cycles $\mathcal{C} = \{C_1, \ldots, C_r\}$ such that none of the vertices in the closed interior of C_r is in T. □

We are finally in the position to state and prove our main theorem:

Theorem 6.2.1. *Let $(G, \mathcal{P}, \tau)$ be an instance for* PD-DP*, where $\mathcal{P} = \{(s_i, t_i) : i \in [k]\}$, for some positive integer k. If* **tw**$(G) \geq 9\sqrt{2k+1} \cdot (k \cdot (2^{2k+2} + 4) + 1) = h$*, then there exists a pair (x, y) of adjoined vertices in G such that if we apply primal-dual contraction (or removal) to them, then $(G, \mathcal{P}, \tau)$ is a YES-instance for* PD-DP *if and only if $(G^{(x,y)}, \mathcal{P}, \tau)$ is a YES-instance for* PD-DP*.*

Proof. Let $T = \{s_1, \ldots, s_k, t_1, \ldots, t_k\} \subseteq V(G)$. We get from Lemma 6.2.3, that G contains an alternating tight sequence of $r = k \cdot (2^{2k+2} + 4)$ concentric cycles $\mathcal{C} = \{C_1, \ldots, C_r\}$ such that all vertices in T are in the open exterior of C_r.

We assume that G contains a pd-linkage whose pattern is $\mathcal{P}$ and let L be a $\mathcal{C}$-cheap one. It is enough to prove that $V(L \cap C_1) = \emptyset$ because we can then choose two adjoined vertices $x, y \in V(C_1) \setminus V_c(G)$ and apply primal-dual contraction to them while retaining the existence of L.

If $k = 1$, thenthe fact that L is $\mathcal{C}$-cheap implies that $L \cap D_{r-1} = \emptyset \Rightarrow V(L \cap C_1) = \emptyset$ and we are done. Therefore, we can assume that $k \geq 2$.

For every $i \in \{1, \ldots, r\}$, we define $\mathcal{Q}^{(i)} = (\mathcal{C}^{(i)}, L^{(i)})$ where $\mathcal{C}^{(i)} = \{C_1, \ldots, C_i\}$ and $L^{(i)}$ is the subgraph of L that is the the union of all connected components of L that intersect D_i. As $r > k$, there exists some $i \in [r]$ such that $\mathcal{Q}^{(i)}$ is touch-free and we let $\mathcal{Q}' = (\mathcal{C}', L') = (\mathcal{C}^{(h)}, L^{(h)})$ where $h = \max\{i \in [r] : \mathcal{Q}^{(i)}$ is a touch-free CL-configuration$\}$. Clearly, $\mathcal{C}'$ is tight in G and L' is $\mathcal{C}'$-cheap. We set $d = r - h$ and observe that the order of the pd-linkage L' is at most $k - d$ while $\mathcal{C}'$ has $r' = r - d > 0$ concentric cycles. Using the same argument as previously, we can assume that $k' \geq 2$, where k' is the order of L'. Additionaly, $k' \leq k - d$, therefore $0 \leq d \leq k - 2$.

As $\mathcal{C}'$ is tight in G and L' is $\mathcal{C}'$-cheap, by Lemma 4.2.2, $\mathcal{Q}'$ is convex. To prove that $V(L \cap C_1) = \emptyset$ it is enough to show that all segments of $\mathcal{Q}$ have positive eccentricity at least 2 which is equivalent to all segments of $\mathcal{Q}'$ having eccentricity at least 2. Assume to the contrary that some segment P_1 of $\mathcal{Q}'$ has eccentricity 1. Then, from the third condition in the definition of convexity we can derive the existence of a sequence $P_1, \ldots, P_{r'-1}$ of segments such that for each $i \in \{1, \ldots, r'\}$, P_{i+1} is inside the zone of P_i.

This in turn implies the existence in $T(\mathcal{Q}')$ of a path of length r' from its root to one of its leaves, therefore $T(\mathcal{Q}')$ has height r'. Then, we get from Lemma 5.2.1 that the real height of $T(\mathcal{Q}')$ is at most $2k' - 3$ and therefore, Observation 5.2.1 gives that the dilation of $T(\mathcal{Q}')$ is at least $\frac{r'}{2k'-3} \geq \frac{k \cdot (2^{2k+2}+4) - d}{2k - 2d} > \frac{k \cdot (2^{2k+2}+4)}{2k} = 2^{2k+1} + 2$. Now, Lemma 5.3.1 gives that G contains an L'-tidy tilted grid $\mathcal{U} = (\mathcal{X}, \mathcal{Z})$ of capacity at least 2^{2k}. Finally, we get from Lemma 6.1.6 that G contains another pd-linkage L'' which is equivalent to L' and such that $c(L'') < c(L')$, which is a contradiction because L' was chosen to be $\mathcal{C}'$-cheap. This concludes our proof. □

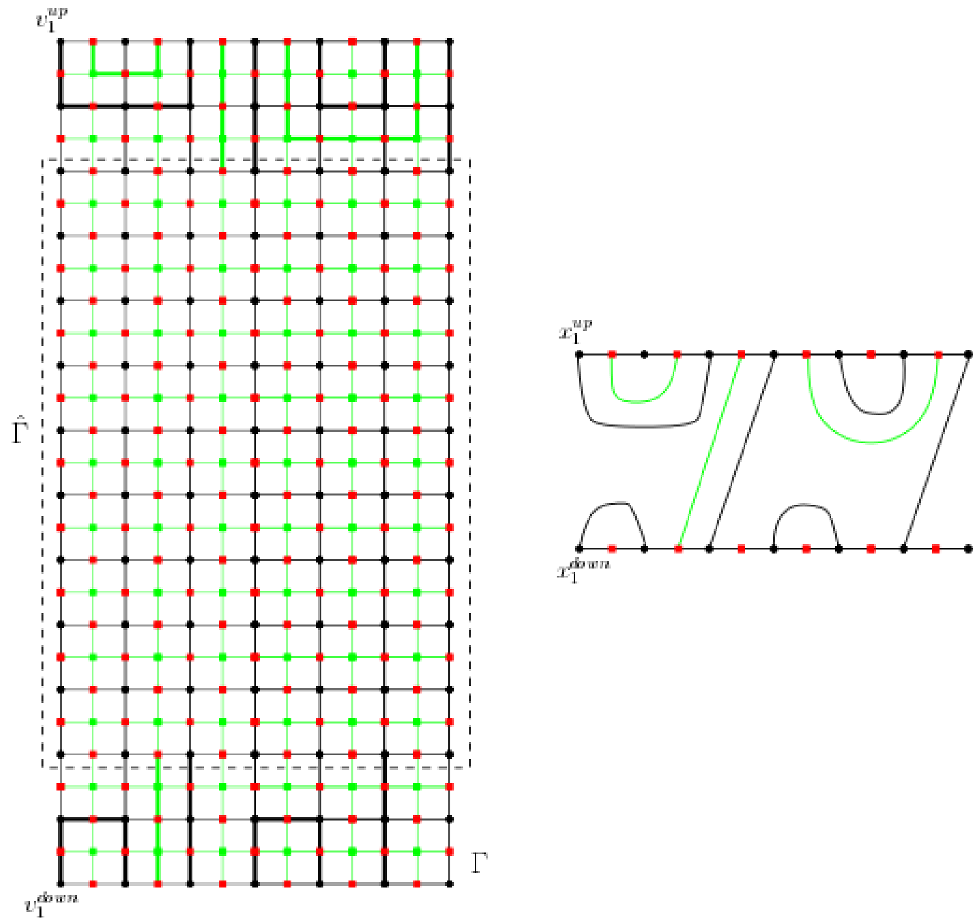

Figure 6.2: An example of the proof of Lemma 6.1.5. At the right: A pd-grid (Γ, χ), where Γ is a (13×26)-pd-grid and the black dots correspond to $\chi^{-1}(p)$, the green squares to $\chi^{-1}(d)$, and the red squares to $\chi^{-1}(c)$. The pd-grid enclosed in the dashed rectangle is a pd-grid $(\hat{\Gamma}, \chi')$, where $\hat{\Gamma}$ (as it is defined in the proof) is a (13×18)-pd-grid and $\chi' = \chi|_{V(\hat{\Gamma})}$.

CHAPTER 7

SKETCH OF AN ALGORITHM AND APPLICATIONS

7.1 Turning the result into an algorithm

We conclude by sketching a way to turn our result into an FPT-algorithm for the PD-DP problem.

The main theorem of this thesis is Theorem 6.2.1 which has a structural essence and states that if the treewidth of our input pd-graph is sufficiently large, then there exists some adjoined pair of vertices (in the central part of some pd-grid) that are in some sense irrelevant for the problem PD-DP. Although not stated explicitly here, this pair can be found algorithmically in time that is polynomial to the size n of the pd-graph.

We claim that given an FPT-algorithm, say $\mathcal{A}$, for solving PD-DP parameterized by the treewidth of the pd-graph in the input, one can construct an FPT-algorithm for PD-DP parameterized by the order of the required linkage. Such an algorithm, with input $\mathcal{I} = (G, \mathcal{P}, \tau)$, works roughly as follows:

Step 1. Check if the treewidth of G is "large" (larger that the function of k that appears in Theorem 6.2.1). If this is the case find an "irrelevant crossing", apply primal-dual contraction and iterate on the new instance. If the answer is no proceed to Step 2.

Step 2. At this point, an equivalent instance $\mathcal{I}' = (G', \mathcal{P}, \tau)$ has been produced, where the treewidth of G' is bounded by a function of k. Employ algorithm $\mathcal{A}$ to solve the problem on $\mathcal{I}'$.

The "traditional" method for solving such kind of problems in graphs of bounded treewidth is dynamic programming or expressing the problem in MSOL (Monadic Second Order Logic) and invoking the celebrated (meta)theorem of Courcelle. Before stating it we briefly describe MSOL.

Monadic Second Order Logic. The syntax of *Monadic Second Order Logic* (MSOL) requires an infinite number of individual variables (we usually use letters

$x, y, z \ldots$) and infinite number of set variables (we usually use capital letters $X, Y, Z, \ldots$). Monadic second-order formulas in the language of graphs are built up from

- atomic formulas $E(x, y)$ (adjacency between x and y), $x = y$ (equality), $X(\mathbf{x})$ (for some set variable X and individual variables **x** and means that vertex x is in set X) by using negation ($\neg$), conjunction ($\wedge$), disjunction ($\vee$), logical implication ($\rightarrow$), and bi-implication ($\leftrightarrow$).

- existential quantification $\exists x$, $\exists X$ and universal quantification $\forall x$, $\forall X$ over individual variables and set variables.

The semantics of MSOL are defined in the obvious way. Let π be a graph property, G be a graph and $\mathcal{G}$ be the class of all graphs. If G has property π we write $\pi(G)$. We say that π is expressible in MSOL if there exists some MSOL formula ϕ_π such that

$$(\forall G \in \mathcal{G})\, [\, G \models \phi_\pi \;\Leftrightarrow\; \pi(G)\,]$$

Our statement of Courcelle's theorem, that follows, lacks formality but for more details on MSOL and Courcelle's theorem we refer the reader to Chapter 13 of [10].

Proposition 7.1.1 ([5]). *If π is a graph property that is expressible in* MSOL*, then there exists some computable function f such that π can be decided in linear (to the size of the graph) time on the class of graphs of bounded treewidth.*

In our case we face the problem of graphs being treated as topological structures (embedded graphs) and not just as combinatorial ones. A way to overcome this difficulty is to "encode" the topology of our plane graph G into another graph G' which is uniquely embeddable (3-connected for example), and transform the property π that we want to decide into another property π' such that π' is expressible in MSOL (hence, from Courcelle's theorem, we automatically have a linear time algorithm for graphs of bounded treewidth) and $\pi(G)$ if and only if $\pi'(G')$. In order to achieve this we can try working an even more *enhanced* version of the pd-graphs, either by considering the pd-graph of the pd-graph of our plane graph G or using other modifications like the *radial graph* (for a more detailed presentation of the previous idea, we refer the reader to [4]). In any case, the goal is to increase the connectivity of our initial plane graph and at some point force it to be *uniquely embeddable* and thus being able to use our combinatorial toolkit.

7.2 Applications

In this section we describe some problems for which our approach can be proved useful. For this we define a notion of critical minor containment.

Plane minors. Given two plane graphs J and H we say that J *critically contains* H *as a plane minor*, and we denote this by $H \leq^{\bullet}_{\mathsf{mn}} J$, if there exists some injection $\phi : V(J) \to V(H)$ such that

1. for every $e = \{x, y\} \in E(H)$, there exists a *unique* edge e' in J with one endpoint in $\phi^{-1}(x)$, and the other in $\phi^{-1}(y)$. We call each such edge e' of J *bridge edge* of J. We also say that e is the *counterpart* of e' in H and that e' is the *counterpart* of e in J.

2. for every $v \in V(H)$, $J[\phi^{-1}(v)]$ is a tree T_v and each leaf of this tree is incident to an bridge edge.

Notice that if $H \leq^{\bullet}_{\mathsf{mn}} J$, then there is a bijection $\lambda : F(H) \to F(J)$ such that, for every $f \in F(H)$, the counterparts of the edges incident to f are exactly the bridge edges of J that are incident to $\lambda(f)$. We call this bijection *face correspondence* between H and J.

Let G be a plane graph and let J be one of the subgraphs of G such that $H \leq^{\bullet}_{\mathsf{mn}} J$ (observe that there might be several such subgraphs of G). We denote by λ_J the *face correspondence* between H and J.

We say that H is a *plane minor* of G, and we denote this by $H \leq_{\mathsf{mn}} G$, if there exists a subgraph J of G such that $H \leq^{\bullet}_{\mathsf{mn}} J$.

We define two problems, which can be argued (as we see soon) to be quite general:

MULTIWAY PLANE FACIAL SEPARATORS
Input: A plane graph G, a plane graph H where $\delta(H) \geq 2$, a set $T \subseteq V(G)$, and a function $\beta : T \to F(H)$.
Parameter: $k = |T| + |V(H)|$.
Question: Is there a subgraph J of G such that $H \leq^{\bullet}_{\mathsf{mn}} J$ and for every $f \in F(H)$, $\beta^{-1}(f) = T \cap \lambda_J(f)$?

Let G be a plane graph, $T \subseteq V(G)$. A cycle C of G is called T-*avoiding* if $V(C) \cap T = \emptyset$. Given a T-avoiding cycle C bounding the open disks D_1 and D_2, we define $\mathcal{P}(C) = \{T \cap D_1, T \cap D_2\}$, i.e., $\mathcal{P}_T(C)$ is a bipartition of T.

PARTITIONING BY CYCLE SEPARATORS
Input: A plane graph G, a set $T \subseteq V(G)$ with $|T| = r$, and a collection $\mathfrak{P} = \{\mathcal{P}_1, \ldots, \mathcal{P}_q\}$ of bipartitions of T.
Parameter: $k = r + q$.
Question: Is there a collection $\mathcal{C} = \{C_1, \ldots, C_q\}$ of vertex-disjoint T-avoiding cycles such that for each $i \in [q]$, $\mathcal{P}_i = \mathcal{P}_T(C_i)$.

The following two problems are special cases of PARTITIONING BY CYCLE SEPARATORS:

CYCLIC MULTIWAY CUT
Input: A plane graph G, and a collection $\mathcal{C} = \{T_1, \ldots, T_r\}$ of (possibly empty) sets of terminals.
Parameter: $k = |\bigcup \mathcal{C}|$.
Question: Is there a collection of r open disks $D_1, \ldots, D_r$ whose boundaries are disjoint cycles of G and where $T_i = T \cap D_i$, for each $i \in [r]$?

MULTI-CYCLIC SEPARATOR
Input: A plane graph G, two disjoint sets of terminals T_1 and T_2, and an integer r.
Parameter: $k = |T_1| + |T_2|$.
Question: Is there a collection of cycles $C_1, \cdots, C_r$ in G, such that, for each C_i, if D_i^1 and D_i^2 are the open disks bounded by C_i, then $T_1 = T \cap D_i^1$, and $T_2 = T \cap D_i^2$, where $T = T_1 \cup T_2$?

All the previously defined problems share a common characteristic: They ask for "vertices being inside cycles of faces" and as one can observe the input graphs are always plane graphs in order for the notion of *enclosure* into a cycle or into a face to be meaningful (which is not the case for graphs as combinatorial structures, which are just set and a binary relation on it and have nothing to do with subsets of the plane). With some effort, which is not presented here, all the questions for this problems can be translated to finding pd-linkages in the pd-graphs that correspond to the input graphs and this is where our work can contribute to.

CHAPTER 8 CONCLUSION

In this thesis we tried to shed some light, from the point of view of Computer Science, to the area of plane graphs. We defined the notions of pd-graph and pd-linkage and proved some structural results in this context, based on the techniques of [2]. In the previous chapter we illustrated some "fertile ground" where we believe that our results can be useful. It would also be interesting to extend existing algorithmic results in the context of embedded graphs and we think that the class of pd-graphs can be, for several cases, the in between step. One first good candidate would be to construct an fpt-algorithm for checking if a plane graph is a topological minor (or a minor) of another plane graph, while respecting their topology. To conclude, it seems like the area of plane graphs is unexplored, compared to the work that has been done in general in the area of graph algorithms, and we think that is an area worth studying.

BIBLIOGRAPHY

[1] Isolde Adler, Stavros G. Kolliopoulos, Dimitrios M. Thilikos, *Planar Disjoint-Paths Completion*, Algorithmica, pages 1-25, 2015

[2] Isolde Adler, Stavros G. Kolliopoulos, Philipp Klaus Krause, Daniel Lokshtanov, Saket Saurabh, Dimitrios M. Thilikos, *Irrelevant Vertices for the Planar Disjoint Paths Problem*, http://arxiv.org/abs/1310.2378, 2015

[3] Sanjeev Arora, Boaz Barak, *Computational Complexity : A Modern Approach*, Cambridge University Press, 2009

[4] Dimitris Chatzidimitriou, Archontia C. Giannopoulou, Spyridon Maniatis, Clément Requilé, Dimitrios M. Thilikos, Dimitris Zoros, FPT *Algorithms for the Plane Subgraph Completion and the Plane Topological Minor Completion Problem*, manuscript

[5] Bruno Courcelle, *The monadic second-order logic of graphs. I. Recognizable sets of finite graphs*, Information and Computation, Volume 85, Issue 1, pages 12-75, 1990

[6] Marek Cygan, Fedor V. Fomin, Dániel Marx, Saket Saurabh, Lukasz Kowalik, Daniel Lokshtanov, Marcin Pilipczuk, *Parameterized Algorithms*, Springer International Publishing AG, 2015

[7] Anuj Dawar, Martin Grohe, Stephan Kreutzer, *Locally excluding a minor*, Logic in Computer Science (LICS'07), pages 270-279, IEEE Computer Society, 2007

[8] Anuj Dawar, Stephan Kreutzer, *Domination problems in nowhere-dense classes*, IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2009), pages 157–168, 2009.

[9] Reinhard Diestel, *Graph Theory*, GTM 173, 4th edition 2010, Springer-Verlag, Heidelberg

[10] RODNEY G. DOWNEY, MICHAEL R. FELLOWS, *Fundamentals of Parameterized Complexity*, Springer-Verlag, London, 2013

[11] JÖRG FLUM, MARTIN GROHE, *Parameterized Complexity Theory*, Springer-Verlag, Berlin, 2006

[12] MICHAEL R. GAREY, DAVID S. JOHNSON, *Computers and Intractability : A Guide to the Theory of* NP*-completeness*, W.H.Freeman and Co Ltd, 1979

[13] PETR A. GOLOVACH, MARCIN. KAMIŃSKI, DANIEL PAULUSMA, DIMITRIOS M. THILIKOS, *Induced packing of odd cycles in a planar graph*, 20th International Symposium on Algorithms and Computation (ISAAC 2009), Volume 5878 of LNCS, pages 514–523, Springer, Berlin, 2009.

[14] PETR A. GOLOVACH, MARCIN. KAMIŃSKI, SPYRIDON MANIATIS, DIMITRIOS M. THILIKOS, *The parameterized complexity of graph cyclability*, European Symposium of Algorithms (ESA 2014), Volume 8737 of LNCS, pages 492-504, Springer, Berlin, 2014

[15] QIAN-PING GU, HISAO TAMAKI, *Improved bounds on the planar branchwidth with respect to the largest grid minor size*, Algorithmica, Volume 64, Issue 3, pages 416–453, 2012.

[16] RICHARD M. KARP, *On the computational complexity of combinatorial problems* Networks, Volume 5, pages 45-68, 1975

[17] KEN-ICHI KAWARABAYASHI, YUSUKE KOBAYASHI, *The induced disjoint paths problem*, 13th Conference on Integer Programming and Combinatorial Optimization (IPCO 2008), volume 5035 of LNCS, pages 47–61, Springer, Berlin, 2008.

[18] KEN-ICHI KAWARABAYASHI, YUSUKE KOBAYASHI, *Algorithms for finding an induced cycle in planar graphs and bounded genus graphs*, 20th ACM-SIAM Symposium on Discrete Algorithms (SODA 2009), pages 1146–1155, ACM-SIAM 2009.

[19] KEN-ICHI KAWARABAYASHI, BRUCE REED, *Odd cycle packing*, 42nd ACM Symposium on Theory of Computing (STOC 2010) pages 695–704, ACM 2010.

[20] KEN-ICHI KAWARABAYASHI PAUL WOLLAN, *A shorter proof of the graph minor algorithm: the unique linkage theorem* 42nd ACM Symposium on Theory of Computing, STOC 2010, pages 687-694. ACM, 2010.

[21] ATHANASSIOS KOUTSONAS, DIMITRIOS M. THILIKOS, *Planar Feedback Vertex Set and Face Cover: Combinatorial Bounds and Subexponential Algorithms*, Algorithmica, Volume 60, Issue 4, pages 987-1003, 2011

[22] MARK R. KRAMER, JAN VAN LEEUWEN, *The complexity of wire-routing and finding minimum area layouts for arbitrary VLSI circuits*, Advances in Computer Research, Issue 2, pages 129–146, 1984

[23] HARRY R. LEWIS, CHRISTOS H. PAPADIMITRIOU, *Elements of the Theory of Computation*, 2nd edition 1997, Pearson Education (US)

[24] CHRISTOS H. PAPADIMITRIOU, *Computational Complexity*, Addison-Wesley, 1994

[25] MATTHIAS MIDDENDORF, FRANK PFEIFFER, *On the complexity of the disjoint paths problem*, Combinatorica, Volume 13, Issue 1, pages 97–107, 1993.

[26] MICHAEL MITZENMACHER, ELI UPFAL, *Probability and Computing : Randomized Algorithms and Probabilistic Analysis*, Cambridge University Press, 2005

[27] RAJEEV MOTWANI, PRABHAKAR RAGHAVAN, *Randomized Algorithms*, Cambridge University Press, 1995

[28] MORITZ MÜLLER, *Randomized Approximations of Parameterized Counting Problems*, Parameterized and Exact Computation, Volume 4169 of the series LNSC, pages 50-59, Springer, Berlin, 2006

[29] ROLF NIEDERMEIER, *Invitation to Fixed-Parameter Algorithms*, Oxford Lecture Series in Mathematics and Its Applications 31, OUP Oxford, 2006 CHRISTOS H. PAPADIMITRIOU, *Computational Complexity*, Addison-Wesley, 1994

[30] NEIL ROBERTSON, PAUL D. SEYMOUR, *Graph Minors.X. Obstructions to Tree-decomposition*, Journal of Combinatorial Theory, Series B, Volume 52, Issue 2, pages 65-110, 1995

[31] NEIL ROBERTSON, PAUL D.SEYMOUR, *Graph Minors.XIII. The Disjoint Paths Problem*, Journal of Combinatorial Theory, Series B, Volume 63, Issue 1, pages 65-110, 1995

[32] NEIL ROBERTSON, PAUL D. SEYMOUR, *Graph Minors.XX. Wagner's conjecture*, Journal of Combinatorial Theory, Series B, Volume 92, Issue 2, pages 325-357, 2004

[33] NEIL ROBERTSON, PAUL D. SEYMOUR, *Graph Minors.XXI. Graphs with unique linkages*, Journal of Combinatorial Theory, Series B, Volume 99, Issue 3, pages 583-616, 2009

[34] NEIL ROBERTSON, PAUL D. SEYMOUR, *Graph minors.XXII. Irrelevant vertices in linkage problems.*, Journal of Combinatorial Theory, Series B, Volume 102, Issue 2, pages 530-563, 2012

[35] RONITT RUBINFELD, ASAF SHAPIRA, *Sublinear Time Algorithms*, SIAM Journal on Computing, Volume 25, Issue 4, pages 1562–1588, 2011

[36] PAUL D. SEYMOUR, *Disjoint paths in graphs*, Discrete Mathematics, Volume 29, Issue 3, pages 293-309, 1980

[37] YOSSI SHILOACH, *A Polynomial Solution to the Undirected Two Paths Problem*, Journal of the ACM Volume 27 Issue 3, pages 445-456, 1980

[38] Robert Solovay, Volker Strassen, *A Fast Monte-Carlo Test for Primality*, SIAM Journal on Computing, Volume 6, Issue 1, pages 84-84, 1976

[39] Carsten Thomassen, *2-linked graphs*, European Journal of Combinatorics, Volume 1, Issue 4, pages 371-378, 1980

[40] Vijay V. Vazirani, *Approximation Algorithms*, Springer-Verlag Berlin and Heidelberg GmbH and Co. KG, 2001

[41] Jens Vygen, NP-*completeness of some edge-disjoint paths problems*, Discrete Applied Mathematics, Volume 61, Issue 1, pages 83-90, 1995

[42] David P. Williamson, David B. Shmoys, *The Design of Approximation Algorithms*, Cambridge University Press, 2011