

Irrelevant Vertices for the Planar Disjoint Paths Problem^a

Isolde Adler^b Stavros G. Kolliopoulos^{cd} Philipp Klaus Krause^{be}
Daniel Lokshtanov^{fg} Saket Saurabh^{hfg} Dimitrios M. Thilikos^{ijd}

December 17, 2015

Abstract

The DISJOINT PATHS PROBLEM asks, given a graph G and a set of pairs of terminals $(s_1, t_1), \dots, (s_k, t_k)$, whether there is a collection of k pairwise vertex-disjoint paths linking s_i and t_i , for $i = 1, \dots, k$. In their $f(k) \cdot n^3$ algorithm for this problem, Robertson and Seymour introduced the *irrelevant vertex technique* according to which in every instance of treewidth greater than $g(k)$ there is an “irrelevant” vertex whose removal creates an equivalent instance of the problem. This fact is based on the celebrated *Unique Linkage Theorem*, whose – very technical – proof gives a function $g(k)$ that is responsible for an immense parameter dependence in the running time of the algorithm. In this paper we give a new and self-contained proof of this result that strongly exploits the combinatorial properties of planar graphs and achieves $g(k) = O(k^{3/2} \cdot 2^k)$. Our bound is radically better than the bounds known for general graphs.

Keywords: Graph Minors, Treewidth, Disjoint Paths Problem

^aEmails: Isolde Adler: iadler@informatik.uni-frankfurt.de, Stavros Kolliopoulos: sgk@di.uoa.gr, Philipp Klaus Krause: philipp@informatik.uni-frankfurt.de, Daniel Lokshtanov: daniello@ii.uib.no, Saket Saurabh: saket@imsc.res.in, Dimitrios M. Thilikos: sedthilk@thilikos.info.

^bInstitut für Informatik, Goethe-Universität, Frankfurt, Germany.

^cDepartment of Informatics and Telecommunications, National and Kapodistrian University of Athens, Athens, Greece.

^dCo-financed by the European Union (European Social Fund – ESF) and Greek national funds through the Operational Program “Education and Lifelong Learning” of the National Strategic Reference Framework (NSRF) - Research Funding Program: “*Thalis. Investing in knowledge society through the European Social Fund*”.

^eSupported by a fellowship within the FIT-Programme of the German Academic Exchange Service (DAAD) at NII, Tokyo and by DFG-Projekt GalA, grant number AD 411/1-1.

^fDepartment of Informatics, University of Bergen, Norway.

^gSupported by “Rigorous Theory of Preprocessing, ERC Advanced Investigator Grant 267959” and “Parameterized Approximation, ERC Starting Grant 306992”.

^hThe Institute of Mathematical Sciences, CIT Campus, Chennai, India.

ⁱDepartment of Mathematics, National and Kapodistrian University of Athens, Athens, Greece.

^jAIGCo project-team, CNRS, LIRMM, France.

1 Introduction

One of the most studied problems in graph theory is the DISJOINT PATHS PROBLEM (DPP): *Given a graph G and a set \mathcal{P} of k pairs of terminals, $(s_1, t_1), \dots, (s_k, t_k)$, decide whether G contains k vertex-disjoint paths P_1, \dots, P_k where P_i has endpoints s_i and t_i , $i = 1, \dots, k$.* In addition to its numerous applications in areas such as network routing and VLSI layout, this problem has been the catalyst for extensive research in algorithms and combinatorics [27]. DPP is NP-complete, along with its edge-disjoint or directed variants, even when the input graph is planar [16–18, 28]. The celebrated algorithm of Robertson and Seymour solves it however in $f(k) \cdot n^3$ steps, where f is some computable function [22]. This implies that, when we parameterize DPP by the number k of pairs of terminals, the problem is fixed-parameter tractable. The Robertson-Seymour algorithm is the central algorithmic result of the Graph Minors series of papers, one of the deepest and most influential bodies of work in graph theory.

The basis of the algorithm in [22] is the so-called *irrelevant-vertex technique* which can be summarized very roughly as follows. As long as the input graph G violates certain structural conditions, it is possible to find a vertex v that is *solution-irrelevant*: every collection of paths certifying a solution to the problem can be rerouted to an *equivalent* one, that links the same pairs of terminals, but in which the new paths avoid v . One then iteratively removes such irrelevant vertices until the structural conditions are met. By that point the graph has been simplified enough so that the problem can be attacked via dynamic programming.

The following two structural conditions are used by the algorithm in [22]: (i) G excludes a clique, whose size depends on k , as a minor and (ii) G has treewidth bounded by some function of k . When it comes to enforcing Condition (ii), the aim is to prove that in graphs without big clique-minors and with treewidth at least $g(k)$ there is always a solution-irrelevant vertex. This is the most complicated part of the proof and it was postponed until the later papers in the series [23, 24]. The bad news is that the complicated proofs also imply an *immense* parametric dependence, as expressed by the function f , of the running time on the parameter k . This puts the algorithm outside the realm of feasibility even for elementary values of k .

The ideas above were powerful enough to be applicable also to problems outside the context of the Graph Minors series. During the last decade, they have been applied to many other combinatorial problems and now they constitute a basic paradigm in parameterized algorithm design (see, e.g., [6, 7, 9, 12, 13, 15]). However, in most applications, the need for overcoming the high parameter dependence emerging from the structural theorems of the Graph Minors series, especially those in [23, 24], remains imperative. Hence two natural directions of research are: simplify parts of the original proof for the general case or focus on specific graph classes that may admit proofs with better parameter dependence. An important step in the first direction was taken recently by Kawarabayashi and Wollan in [14] who gave an easier and shorter proof of the results

in [23, 24]. While the parameter dependence of the new proof is certainly much better than the previous, immense, function, it is still huge: a rough estimation from [14] gives a lower bound for $g(k)$ of magnitude $2^{2^{\Omega(k)}}$ which in turn implies a lower bound for $f(k)$ of magnitude $2^{2^{2^{\Omega(k)}}}$.

In this paper we offer a solid advance in the second direction, focusing on planar graphs (see also [20, 26] for previous results on planar graphs). We show that, for planar graphs, $g(k)$ is single exponential. In particular we prove the following result.

Theorem 1. *Every instance of DPP consisting of a planar graph G with treewidth at least $82 \cdot k^{3/2} \cdot 2^k$ and k pairs of terminals contains a vertex v such that every solution to DPP can be replaced by an equivalent one whose paths avoid v .*

The proof of Theorem 1 is presented in Section 3 and deviates significantly from those in [14, 23, 24]. It is self-contained and exploits extensively the combinatorics of planar graphs. Given a DPP instance defined on a planar graph G , we prove that if G contains as a subgraph a subdivision of a sufficiently large (exponential in k) grid, whose “perimeter” does not enclose any terminal, then the “central” vertex v of the grid is solution-irrelevant for this instance. It follows that the “area” provided by the grid is big enough so that every solution that uses v can be rerouted to an equivalent one that does not go so deep in the grid and therefore avoids the vertex v .

Combining Theorem 1 with known algorithmic results, it is possible to reduce, in $2^{2^{O(k)}} \cdot n^2$ steps, a planar instance of DPP to an equivalent one whose graph has treewidth $2^{O(k)}$. Then, using standard dynamic programming on tree decompositions, a solution, if one exists, can be found in $2^{2^{O(k)}} \cdot n$ steps. The parametric dependence of this algorithm is a step forward in the study of the parameterized complexity of DPP on planar graphs. This algorithm is abstracted in the following theorem, whose proof is in Section 4.

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

An extended abstract of this work, without any proofs, appeared in [2]. Some of our ideas have proved useful in the recent breakthrough result of Cygan et al. that establishes fixed-parameter tractability for k -disjoint paths on planar directed graphs [5].

2 Basic definitions

Throughout this paper, given a collection of sets \mathcal{C} we denote by \mathbf{UC} the set $\cup_{x \in \mathcal{C}} x$, i.e., the union of all sets in \mathcal{C} .

All graphs that we consider are finite, undirected, and simple. We denote the vertex set of a graph G by $V(G)$ and the edge set by $E(G)$. Every edge is a two-element subset of $V(G)$. A graph H is a *subgraph* of a graph G , denoted by $H \subseteq G$, if $V(H) \subseteq V(G)$ and

$E(H) \subseteq E(G)$. Given two graphs G and H , we define $G \cap H = (V(G) \cap V(H), E(G) \cap E(H))$ and $G \cup H = (V(G) \cup V(H), E(G) \cup E(H))$. Given a $S \subseteq V(G)$, we also denote by $G[S]$ the subgraph of G induced by S .

A *path* in a graph G is a connected acyclic subgraph with at least one vertex whose vertices have degree at most 2. The *length* of a path P is equal to the number of its edges. The *endpoints* of a path P are its vertices of degree 1 (in the trivial case where there is only one endpoint x , we say that the endpoints of P are x and x). An (x, y) -path of G is any path of G whose endpoints are x and y .

A *cycle* of a graph G is a connected subgraph of G whose vertices have degree 2. For graphs G and H the *cartesian product* is the graph whose vertex set is $V(G) \times V(H)$ and whose edge set is $\{(v, v'), (w, w') \mid (\{v, w\} \in E(G) \wedge v' = w') \vee (v = w \wedge \{v', w'\} \in E(H))\}$.

The DISJOINT PATHS problem. The problem that we examine in this paper is the following.

DISJOINT PATHS (DPP)
Input: A graph G , and a collection $\mathcal{P} = \{(s_i, t_i) \in V(G)^2, i \in \{1, \dots, k\}\}$ of pairs of $2k$ terminals of G .
Question: Are there k pairwise vertex-disjoint paths P_1, \dots, P_k in G such that for $i \in \{1, \dots, k\}$, P_i has endpoints s_i and t_i ?

We call the k -pairwise vertex-disjoint paths certifying a YES-instance of DPP a *solution* of DPP for the input (G, \mathcal{P}) . Given an instance (G, \mathcal{P}) of DPP, we say that a non-terminal vertex $v \in V(G)$ is *irrelevant for* (G, \mathcal{P}) , if (G, \mathcal{P}) is a YES-instance if and only if $(G \setminus v, \mathcal{P})$ is a YES-instance. We denote by PDPP the restriction of DPP on instances (G, \mathcal{P}) where G is a planar graph.

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

- i.* For every two distinct vertices x and y of H , $G[\phi(x)]$ and $G[\phi(y)]$ are two vertex-disjoint connected subgraphs of G and
- ii.* for every two adjacent vertex x and y of H , $G[\phi(x) \cup \phi(y)]$ is a connected subgraph of G .

We call the function ϕ *minor model* of H in 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 incident to more than 4 vertices. The *outer cycle* of an embedding of an $(m \times n)$ -grid is the one that is the boundary of its infinite face. We also refer to the *horizontal* and the *vertical lines* of an $(m \times n)$ -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 defined in the obvious way (see Figure 1 for an example).




Figure 1: A drawing of the (6×6) -grid. The four white round vertices are its corners and the four grey square vertices are its centers. The cycle formed by the “fat” edges is the outer cycle.

Plane graphs Whenever we refer to a planar graph G we consider an embedding of G in the plane $\Sigma = \mathbb{R}^2$. To simplify notation, we do not distinguish between a vertex of G and the point of Σ used in the drawing to represent the vertex or between an edge

and the arc representing it. We also consider a plane graph G as the union of the points corresponding to its vertices and edges. That way, edges and faces are considered to be open sets of Σ . Moreover, a subgraph H of G can be seen as a graph H , where the points corresponding to H are a subset of the points corresponding to G .

Recall that $\Delta \subseteq \Sigma$ is an open (resp. closed) disc if it is homeomorphic to $\{(x, y) : x^2 + y^2 < 1\}$ (resp. $\{(x, y) : x^2 + y^2 \leq 1\}$). Given a cycle C of G we define its *open-interior* (resp. *open-exterior*) as the connected component of $\Sigma \setminus C$ that is disjoint from (resp. contains) the infinite face of G . The *closed-interior* (resp. *closed-exterior*) of C is the closure of its *open-interior* (resp. *open-exterior*). Given a set $A \subseteq \Sigma$, we denote its *interior* (resp. *closure*) by $\mathbf{int}(A)$ (resp. $\mathbf{clos}(A)$). An *open (resp. closed) arc* I in \mathbb{R}^2 is any set homeomorphic to the set $\{(x, 0) \mid x \in (0, 1)\}$ (resp. $\{(x, 0) \mid x \in [0, 1]\}$) and the *endpoints* of I are defined in the obvious way. We also define $\mathbf{trim}(I)$ as the set of all points of the arc I except for its endpoints.

Outerplanar graphs. 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. Notice that if the outerplanar graph G is biconnected, then its weak dual is a tree. We call a face of an outerplanar graph *simplicial* if it corresponds to a leaf of the graph's weak dual.

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

The *width* of (T, χ) is defined as $\mathbf{w}(T, \chi) := \max \{ |\chi(t)| - 1 \mid t \in V(T) \}$.

The *tree-width* of G is defined as

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

We need the next proposition that follows directly by combining the main result of [10] and (5.1) from [21].

Proposition 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.*

Our algorithmic results require the following proposition. It follows from the main result of [19] (see also Algorithm (3.3) in [22]). The parametric dependence of k in the running time follows because the algorithm in [19] uses as a subroutine the algorithm in [4] that runs in $2^{k^{O(1)}} \cdot n$ steps.

Proposition 2. *There exists an algorithm that, given an n -vertex graph G and a positive integer k , either outputs a tree decomposition of G of width at most k or outputs a subgraph G' of G with treewidth greater than k and a tree decomposition of G' of width at most $2k$, in $2^{k^{O(1)}} \cdot n$ steps.*

3 Irrelevant vertices in graphs of large treewidth

In this section we prove our main result, namely Theorem 1. We introduce the notion of cheap linkages and explore their structural properties in Subsections 3.1 and 3.4. In Subsection 3.7 we bring together the structural results to show the existence of an irrelevant vertex in a graph of large treewidth.

3.1 Configurations and cheap linkages

In this subsection we introduce some basic definitions on planar graphs that are necessary for our proof.

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 whose internal vertices belong to the open interior of C .

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} . 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 $D_i \subsetneq D \subsetneq D_{i+1}$.

Lemma 1. *There exists an algorithm that given a positive integer r , an n -vertex plane graph G , and a $T \subseteq V(G)$, either outputs a tree decomposition of G of width at most $9 \cdot (r+1) \cdot \lceil \sqrt{|T|+1} \rceil$ or an internally chordless cycle C of G such that there exists a tight sequence of cycles C_0, \dots, C_r in G where*

- $C_0 = C$ and
- all vertices of T are in the open exterior of C_r .

Moreover, this algorithm runs in $2^{(r \cdot \sqrt{|T|})^{O(1)}} \cdot n$ steps.




Figure 2: An example of a plane graph G and a tight sequence of 3 concentric cycles in it. Notice that the addition to G of any of the dashed edges makes this collection of cycles non-tight.

Proof. Let $x = |T| + 1$ and $y = 2(r + 1) \cdot \lceil \sqrt{x} \rceil$. From Proposition 1, if $\mathbf{tw}(G) \geq 4.5 \cdot y + 1$, then G contains as a minor a $(y \times y)$ -grid Γ . We now observe that the grid Γ contains as subgraphs x pairwise disjoint $(2(r + 1) \times 2(r + 1))$ -grids $\Gamma_1, \dots, \Gamma_x$. Note that each $\Gamma_i, i \in \{1, \dots, x\}$ contains a sequence of $r + 1$ concentric cycles that, given a minor model ϕ of Γ in G , can be used to construct, in linear time, a sequence of $r + 1$ concentric cycles $\mathcal{C}_i = \{C_0^i, C_1^i, \dots, C_r^i\}$ in G such that for every $i, j \in \{1, \dots, x\}$, where $i \neq j$, all cycles in \mathcal{C}_j are in the open exterior of C_r^i .

Note that at least one, say C_r^i , of the cycles in $\{C_r^1, \dots, C_r^x\}$ should contain all the vertices of T in its open exterior. Let e be any edge of C_0^i . Let also f be the face of G that is contained in the open interior of C_0^i and is incident to e . Let J_f be the graph consisting of the vertices and the edges that are incident to f . It is easy to verify that, J_f contains an internally chordless cycle C that contains the edge e . Given C_0^i , the cycle C can be found in linear time. Notice now that G contains a tight sequence of cycles C_0, C_1, \dots, C_r such that $C_0 = C$ and where, for $h \in \{0, \dots, r\}$, C_h is in the closed interior of C_h^i . The result follows as the open exterior of C_r contains the open exterior of C_r^i and therefore contains all vertices in T .

The algorithm runs as follows: first uses the algorithm of Proposition 2 for $k = 4.5 \cdot y$. If the algorithm outputs a tree decomposition of G of width at most k , then we are done. Otherwise it outputs a subgraph G' of G where $\mathbf{tw}(G') > k$ and a tree decomposition of G' of width $\leq 2k$. We use this tree decomposition in order to find a minor model ϕ of the $(y \times y)$ -grid Γ in G' . This can be done in $2^{k^{O(1)}} = 2^{(r \cdot \sqrt{|T|})^{O(1)}} \cdot n$ steps using the algorithm in [1] (or, alternatively, the algorithm in [11]). Clearly, ϕ is also a minor

model of Γ in G . We may now use ϕ , as explained above, in order to identify, in linear time, the required internally chordless cycle C in G . \square

Linkages. A *linkage* in a graph G is a non-empty subgraph L of G whose connected components are all paths. The *paths* of a linkage are its connected components and we denote them by $\mathcal{P}(L)$. The *terminals* of a linkage L are the endpoints of the paths in $\mathcal{P}(L)$, and the *pattern* of L is the set $\{\{s, t\} \mid \mathcal{P}(L) \text{ contains a path from } s \text{ to } t \text{ in } G\}$. Two linkages are *equivalent* if they have the same pattern.

Segments. Let G be a plane 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 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 $\mathcal{P}(L)$.




Figure 3: An example of a CL-configuration $\mathcal{Q} = (\mathcal{C}, L)$ where \mathcal{C} contains 5 cycles and L has 7 paths. \mathcal{Q} has 13 segments. Linkage paths $A, B, C, D, E, F,$ and $G,$ contain 2, 2, 2, 1, 1, 2, 3 of these segments respectively. Also the eccentricities of the segments of $A,$ are 0 and 2, of B are 3 and 4. Notice that one of the two segments of A has two 3-chords, each having 2 semi 3-chords.

CL-configurations. Given a plane graph G , we say that a pair $\mathcal{Q} = (\mathcal{C}, L)$ is a *CL-configuration* of G of *depth* r if $\mathcal{C} = \{C_0, \dots, C_r\}$ is a sequence of concentric cycles in G , L is a linkage of G , and D_r does not contain any terminals of L . 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 has eccentricity r , i.e., it is a subpath 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 $P \cap \text{int}(D_i)$ (notice that i -chords are open arcs). 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 arcs). 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).




Figure 4: An example of a CL-configuration (\mathcal{C}, L) where the linkage L is \mathcal{C} -cheap. Only the 5 concentric cycles of \mathcal{C} and a cropped part of the linkage L are depicted. Notice that the collection of concentric cycles \mathcal{C} is not tight.

A CL-configuration $\mathcal{Q} = (\mathcal{C}, L)$ is called *reduced* if the graph $L \cap \mathbf{UC}$ is edgeless. 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 $L \cap \mathbf{UC}$. We then define G^* as the graph obtained if we contract in G all edges in E^\bullet . We also define \mathcal{Q}^* as the pair (\mathcal{C}^*, L^*) obtained if in L and in the cycles of \mathcal{C} we contract all edges of E^\bullet . Notice that \mathcal{Q}^* is a reduced CL-configuration of G^* . We call (\mathcal{Q}^*, G^*) the *reduced pair* of G and \mathcal{Q} .

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

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

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

Observation 1. 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

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




Figure 5: An example of a convex CL-configuration (\mathcal{C}, L) . In the picture, only the 5 cycles in \mathcal{C} and a cropped portion of L is depicted.

3.2 Convex configurations

We introduce CL-configurations with particular characteristics that will be useful for the subsequent proofs. We then show that these characteristics are implied by tightness and cheapness.

Convex CL-configurations. A segment P of \mathcal{Q} is *convex* if the following three conditions are satisfied:

- (i) it has no 0-chord and
- (ii) for every $i \in \{1, \dots, r\}$, the following hold:
 - a. P has at most one i -chord
 - b. if P has an i -chord, then $P \cap C_{i-1} \neq \emptyset$.

- c. Each i -chord of P has exactly two i -semichords.
- (iii) If P has eccentricity $i < r$, there is another segment inside the zone of P with eccentricity $i + 1$.

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

Observation 2. Let $\mathcal{Q} = (C, L)$ be a CL-configuration and let $(G^*, \mathcal{Q}^* = (C^*, L^*))$ be the reduced pair of G and \mathcal{Q} . Then \mathcal{Q} is convex if and only if \mathcal{Q}^* is convex.




Figure 6: A visualization of the conditions of Lemma 2.

The proof of the following lemma uses elementary topological arguments.

Lemma 2. Let Δ_1, Δ_2 be closed disks of \mathbb{R}^2 where $\mathbf{int}(\Delta_1) \cap \mathbf{int}(\Delta_2) = \emptyset$ and such that $\Delta_1 \cup \Delta_2$ is also a closed disk. Let $\Delta_3 = \mathbb{R}^2 \setminus \mathbf{int}(\Delta_1 \cup \Delta_2)$ and let $Y = \mathbf{bnd}(\Delta_3) \cap \Delta_2$ and $Q = \mathbf{trim}(\Delta_1 \cap \Delta_2)$. Let 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 $\mathbf{int}(\Delta_1) \cap P$ 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 $\mathbf{int}(\Delta_1)$ 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 $Q' \subseteq \mathbf{int}(\Delta_1)$, $\mathbf{int}(\Delta_1) \cap P$ has at least one connected component.

Assume now that $\mathbf{int}(\Delta_1) \cap P$ 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 $\mathbf{int}(\Delta_1) \cap P$ has exactly one connected component, it holds that $P' \cap \mathbf{int}(\Delta_1) = \emptyset$. Let Q' be an open arc that is a subset of $\mathbf{int}(\Delta_1)$ 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 $\mathbf{int}(\Delta_1)$ or Y , a contradiction as $P' \subseteq P$ and $Y \cap P = \emptyset$. \square

Lemma 3. *Let G be a plane graph and $\mathcal{Q} = (\mathcal{C}, L)$ 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. By Observations 1 and 2, we may assume that \mathcal{Q} is reduced. Consider any segment of \mathcal{Q} . We show that it satisfies the three conditions of convexity. Conditions (i) and (ii).b follow directly from the tightness of \mathcal{C} . Condition (iii) follows from the fact that L is \mathcal{C} -cheap. In the rest of the proof we show Conditions (ii).a. and (ii).c. For this, we consider the minimum $i \in \{0, \dots, r\}$ such that one of these two conditions is violated. From Condition (i), $i \geq 1$. Let W be a segment of \mathcal{Q} containing an i -chord X for which one of Conditions (ii).a, (ii).c is violated.

We now define the set Q according to which of the two conditions is violated. We distinguish two cases:

Case 1. Condition (ii).c is violated. From Condition (ii).b, $X \setminus D_{i-1}$ contains more than two i -semichords of X . Let J_1 be the biconnected outerplanar graph defined by the union of C_{i-1} and the i -semichords of X that do not have an endpoint on C_i . As there are at least three i -semichords in X , J_1 has at least one internal edge and therefore at least two simplicial faces. Moreover there are exactly two i -semichords of X , say K_1 , K_2 , that have an endpoint in C_i and K_1 and K_2 belong to the same, say F' , face of J_1 . Let Δ_2 be the closure of a simplicial face of J_1 that is not F' .




Figure 7: The two cases of the proof of Lemma 3. In the left part is depicted an i -chord X that has 6 i -semichords and in the right part is depicted a segment W and the way it crosses the cycles D_i and D_{i-1} . In the figure on the right the segment W has 7 i -chords and 14 i -semichords.

Case 2. Condition (ii).c holds while Condition (ii).a is violated. Let J_2 be the bicon-

nected outerplanar graph defined by the union of C_{i-1} and the connected components of $W \setminus D_{i-1}$ that do not contain endpoints of W in their boundary. Notice that the rest of the connected components of $W \setminus D_{i-1}$ are exactly two, say K_1 and K_2 . Notice that K_1 and K_2 are subsets of the same face, say F' , of J_2 . As there are at least two i -chords in W , J_2 contains at least one internal edge and therefore at least two simplicial faces. Let Δ_2 be the closure of a simplicial face of J_2 that is not F' .

In both of the above cases, we set $\Delta_1 = D_{i-1}$, $\Delta_3 = \mathbb{R}^2 \setminus \mathbf{int}(\Delta_1 \cup \Delta_2)$, $Y = \mathbf{bnd}(\Delta_3) \cap \Delta_2$, and $Q = \mathbf{trim}(\Delta_1 \cap \Delta_2)$. Notice that $Y = \mathbf{bnd}(\Delta_2) \setminus Q$, therefore $Y \subseteq W$.

We claim that $L \cap Q \neq \emptyset$. Suppose not. We consider W' as the path in $W \cup Q$ that contains Q as a subset and has the same endpoints as W . Then, $L' = (L \setminus W) \cup W'$ is a linkage, equivalent to L , where $c(L') < c(L)$, a contradiction to the fact that L is \mathcal{C} -cheap. We just proved that $L \cap Q \neq \emptyset$ which in turn implies that L contains a segment P for which $P \cap Q \neq \emptyset$. We distinguish two cases:

Case A. $W \neq P$. This implies that $W \cap P = \emptyset$. As $Y \subseteq W$, it follows that $Y \cap P = \emptyset$. Therefore, by Lemma 2, $\mathbf{int}(\Delta_1) \cap P$ has at least two connected components, therefore P has at least two $(i-1)$ -chords. If $i > 1$, then Condition (ii).a is violated for $i-1$, which contradicts the choice of i . If $i = 1$, then P has at least one 0-chord, which violates Condition (i), that, as explained at the beginning of the proof, holds for every segment of Q .

Case B. $W = P$. Recall that $Y \subseteq W$ therefore $Y \subseteq P$. Let p_1 and p_2 be the endpoints of Q . As Q is reduced 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

- p_i is an endpoint of $Z_i, i \in \{1, 2\}$.
- $Z_i \subseteq \mathbf{clos}(Q), i \in \{1, 2\}$, and
- $P \cap Z_i = \{p_i\}, i \in \{1, 2\}$.

Consider also a closed arc Y' that is a subset of $\mathbf{int}(\Delta_2) \cup \{p'_1, p'_2\}$ that does not intersect L and whose endpoints are p'_1 and p'_2 . Let now $\Delta'_1 = \Delta_1$, let Δ'_2 be the closed disk defined by the cycle $\mathbf{clos}(Q \setminus (Z_1 \cup Z_2)) \cup Y'$ that is a subset of Δ_2 . Let also $\Delta'_3 = \mathbb{R}^2 \setminus \mathbf{int}(\Delta'_1 \cup \Delta'_2)$ and $Q' = \mathbf{trim}(\Delta'_1 \cap \Delta'_2)$. 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 2, $\mathbf{int}(\Delta'_1) \cap P$ has at least two connected components. Therefore P has at least two $(i-1)$ -chords. This yields a contradiction, as in Case A. \square

3.3 Bounding the number of extremal segments

In this subsection we prove that the number of extremal segments is bounded by a linear function of the number of linkage paths.

Out-segments, hairs, and flying hairs. Let G be a plane 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 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 linkage L of G and a closed disk D of \mathbb{R}^2 , we define $\mathbf{out}_D(L)$ to be the graph obtained from $(L \cup \mathbf{bnd}(D)) \setminus \mathbf{int}(D)$ after dissolving all vertices that do not belong in L . For example $\mathbf{out}_{D_r}(L)$ 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 . Let f be a face of $\mathbf{out}_{D_r}(L)$ that is different from $\mathbf{int}(D_r)$. We say that f is a *cave of $\mathbf{out}_{D_r}(L)$* if the union of the out-segments and extremal segments in the boundary of f is a connected set. Recall that a segment of \mathcal{Q} is extremal if it 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_2 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 exists a sequence of edges $e_1, \dots, e_r \in E$ where $e_1 = e$, $e_r = e'$ and for each $i \in \{1, \dots, r-1\}$ 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 of L , the number of the connected components of $P \cap C_r$ is not 1.

Lemma 4. *Let G be a plane 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. Let $(G^*, \mathcal{Q}^* = (\mathcal{C}^*, L^*))$ be the reduced pair of G and \mathcal{Q} . Notice that, by Observation 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 all trivial paths (i.e., paths consisting of only one vertex). 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. 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 $\mathbf{bnd}(f)$ all the points that belong to out-segments. Clearly, R results from a subpath R^+ of C_r^* after removing its endpoints, i.e., $R = \mathbf{trim}(R^+)$.

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 $L^{*'} = (L^* \setminus (\mathbf{bnd}(f))) \cup R$ is also a linkage with the same pattern as L^* where $c(L^{*'}) < c(L^*)$, 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 3, \mathcal{Q}^* is convex, therefore 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 L^* lie outside D_r^* , the set $P' = P \cap (\mathbf{bnd}(f) \setminus D_r^*)$ 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 minus the number of bouncing terminals. Therefore,

$$\rho = |T_2| + |S|. \quad (1)$$

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

$$|S| \leq |T_1| - 2 \quad (2)$$




Figure 8: Examples of the graphs J and J^- in the proof of Lemma 4 (the outer face in the picture corresponds to the interior of D_r). The faces that are caves contain the word cave. FH: flying hair, BH: bouncing hair, IH: invading hair. The forest $\Psi = J^- \setminus E(C_r)$ has 6 edges and 4 connected components. The weak dual T of J^- is depicted with dashed lines. The large white square vertices are the rich vertices of T .

Indeed the truth of (2) along with (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 (2). For this, we need two more claims.

Claim 2: $|F_0| \leq \kappa(\Psi) - 1$.

Proof. We use induction on $\kappa(\Psi)$. Let $K_1, \dots, K_{\kappa(\Psi)}$ 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 $E(K_i)$, $i \in \{1, \dots, \kappa(\Psi)\}$, and observe that $E(K_1^*), \dots, E(K_{\kappa(\Psi)}^*)$ 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 $\{K_1^*, \dots, K_{\kappa(\Psi)}^*\}$, otherwise it is called *poor* (see Figure 8). 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 $\{K_1, \dots, K_{\kappa(\Psi)}\}$ is equivalent to the existence of a peripheral subtree in $\{K_1^*, \dots, K_{\kappa(\Psi)}^*\}$. 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 $\{K_1^*, \dots, K_{\kappa(\Psi)}^*\}$. 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, (2) holds and this completes the proof of the lemma. \square

3.4 Bounding the number and size of segment types

In this section we introduce the notion of segment type that partitions the segments into classes of mutually “parallel” segments. We next prove that, in the light of the results of the previous section, the number of these classes is bounded by a linear function of the number k of linkage paths. In Subsections 3.5 and 3.5 we show that if one of these equivalence classes has size more than 2^k , then an equivalent cheaper linkage can be found. All these facts will be employed in the culminating Subsection 3.7 in order to prove that a cheap linkage cannot go very “deep” into the cycles of a cheap CL-configuration. That way we will be able to quantify the depth at which an irrelevant vertex is guaranteed to exist.

Types of segments. Let G be a plane 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 the two paths on C_r connecting an endpoint of S_1 with an endpoint of S_2 and passing 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 *type of segment* is an equivalence class of segments of \mathcal{Q} under the relation \parallel .

Given a 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 is the dual of the 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, \dots, W_{ρ_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 $i \in \{1, \dots, \rho_l\}$, 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.

Observation 3. *Let G be a plane graph and let $\mathcal{Q} = (\mathcal{C}, L)$ be a convex CL-configuration of G . Then the dilation of $T(\mathcal{Q})$ is equal to the maximum cardinality of an equivalence class of \parallel .*

Observation 4. *Let G be a plane 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 following lemma is an immediate consequence of Lemma 4 and the definition of a segment tree. The condition that $L \cap C_r \neq \emptyset$ simply requires that the CL-configuration that we consider is non-trivial in the sense that the linkage L enters the closed disk D_r .

Lemma 5. *Let G be a plane 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. Certainly, the convexity of \mathcal{Q} follows directly from Lemma 3. 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 segment all of which are connected components of $C_r \cap P$. Let P_1 and P_2 be the closures




Figure 9: The graph $\mathbf{in}_{D_r}(L)$ for some convex CL-configuration $\mathcal{Q} = (\mathcal{C}, L)$ and the tree $T(\mathcal{Q})$. Internal edges in $\mathbf{in}_{D_r}(L)$ of the same type are drawn as lines of the same type. \mathcal{Q} has 11 extremal segments, as many as the leaves of $T(\mathcal{Q})$. The relation \parallel has 19 equivalent classes. The dilation of $T(\mathcal{Q})$ is 4, its height is 8 and its real height is 4.

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 $P_1 \cup P' \cup P_2$ is a cheaper linkage with the same pattern as L , 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. This completes the proof that $|\mathcal{P}(L)| \geq 2$.

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 4, $T(\mathcal{Q})$ 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 $2 \cdot |\mathcal{P}(L)| - 3$. \square

3.5 Tidy grids in convex configurations

In this subsection we prove that the existence of many “parallel” segments implies the existence of a big enough grid-like structure.

Topological minors. We say that a graph H is a *topological minor* of a graph G if there exists an injective function $\phi_0 : V(H) \rightarrow V(G)$ and a function ϕ_1 mapping the edges of H to paths of 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)* .

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

- for each $i, j \in \{1, \dots, r\}$ $I_{i,j} = X_i \cap Z_j$ is a (possibly edgeless) path of G ,
- for $i \in \{1, \dots, r\}$ the subpaths $I_{i,1}, I_{i,2}, \dots, I_{i,r}$ appear in this order in X_i .
- for $j \in \{1, \dots, r\}$ the subpaths $I_{1,j}, I_{2,j}, \dots, I_{r,j}$ 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

$$G_{\mathcal{U}} = \left(\bigcup_{i \in \{1, \dots, r\}} X_i \right) \cup \left(\bigcup_{i \in \{1, \dots, r\}} Z_i \right)$$

and let $G_{\mathcal{U}}^*$ be the graph taken from the graph after contracting all edges in $\bigcup_{(i,j) \in \{1, \dots, r\}^2} I_{i,j}$. Then $G_{\mathcal{U}}^*$ contains the $(r \times r)$ -grid Γ as a topological minor via a pair (χ_0, χ_1) such that

- the upper left (resp. upper right, down right, down left) corner of Γ is mapped via χ_0 to the (single) endpoint of $I_{1,1}$ (resp. $I_{1,r}, I_{r,r}$, and $I_{r,1}$).
- $\bigcup_{e \in E(\Gamma)} \chi_1(e) = G_{\mathcal{U}}^*$ (this makes $G_{\mathcal{U}}^*$ to be a subdivision of Γ).

We call the subgraph $G_{\mathcal{U}}$ of G *realization* of the tilted grid \mathcal{U} and the graph $G_{\mathcal{U}}^*$ *representation* of \mathcal{U} . We treat both $G_{\mathcal{U}}$ and $G_{\mathcal{U}}^*$ as plane graphs. We also 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 cycle $X_1 \cup Z_1 \cup X_r \cup Z_r$. Given a graph G and a linkage L of G we say that a tilted grid $\mathcal{U} = (\mathcal{X}, \mathcal{Z})$ of G is an *L -tidy tilted grid* of G if $D_{\mathcal{U}} \cap L = \mathbf{UZ}$ where $D_{\mathcal{U}}$ is the closed-interior of the perimeter of $G_{\mathcal{U}}$.

Lemma 6. *Let G be a plane 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 \parallel . Then G contains a tilted grid $\mathcal{U} = (\mathcal{X}, \mathcal{Z})$ of capacity $\lceil |\mathcal{S}|/2 \rceil$ that is an L -tidy tilted grid of G .*

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

- $\sigma_1 = \sigma^{\min}$,




Figure 10: A visualisation of the proof of Lemma 6. It holds that $\sigma_1 = 0$ and $\sigma_5 = 4$, $m = 5$, $m' = 3$. The two shadowed regions indicate the two connected components of $D_S \cap A_C$.

- $\sigma_m = \sigma^{\max}$, and
- for all $i \in \{1, \dots, m-1\}$, $\sigma_{i+1} = \sigma_i + 1$.
- for all $i \in \{1, \dots, m\}$, $I_{i, \sigma_i} = S_i \cap C_{\sigma_i}$ is a 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 closed annulus defined by 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

$$(L \cup \mathbf{UC}) \cap \Delta.$$

It is now easy to verify that the above graph is the realization $G_{\mathcal{U}}$ of a tilted grid $\mathcal{U} = (\mathcal{X}, \mathcal{Z})$ of capacity m' , where the paths in \mathcal{X} are the portions of the cycles

$C_{\sigma^{\max}-(m'-1), \dots, C_{\sigma^{\max}}$ cropped by Δ , while the paths in \mathcal{Z} are the portions of the paths in $\{S_1, \dots, S_{m'}\}$ cropped by Δ (see Figure 10). As \mathcal{S} is an equivalence class of \parallel , it follows that \mathcal{U} is L -tidy, as required. \square

3.6 Replacing linkages by cheaper ones

In this section we prove that a linkage L of k paths can be rerouted to a cheaper one, given the existence of an L -tidy tilted grid of capacity greater than 2^k . Given that L is a cheap linkage, this will imply an exponential upper bound on the capacity of an L -tidy tilted grid.

Let G be a plane graph and let L be a 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. This naturally partitions the vertices of $\mathbf{bnd}(D) \cap L$ into the *up* and *down* ones. The following proposition is implicit in the proof of Theorem 2 in [3] (see the derivation of the unique claim in the proof of the former theorem). See also [8] for related results.

Proposition 3. *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 $\mathbf{bnd}(D) \cap L$, such that there exists some linkage R that is a subgraph of $L \setminus \mathbf{int}(D)$ such that $R \cup \mathbf{UN}$ is a linkage of the graph $(G \setminus D) \cup \mathbf{UN}$ that is equivalent to L .*




Figure 11: An example of the proof of Lemma 7, where $k = 16$, $k' = 11$, and $\rho = 5$. The white vertices of the higher (resp. lower) horizontal line are the vertices in $\{p_1^{\text{up}}, \dots, p_5^{\text{up}}\}$ (resp. $\{p_1^{\text{down}}, \dots, p_5^{\text{down}}\}$).

Lemma 7. *Let k, k', ρ be integers such that $0 \leq \rho \leq k' \leq k$. Let Γ be a $(k \times k')$ -grid and let $\{p_1^{\text{up}}, \dots, p_\rho^{\text{up}}\}$ (resp. $\{p_1^{\text{down}}, \dots, p_\rho^{\text{down}}\}$) be vertices of the higher (resp. lower) horizontal line arranged as they appear in it from left to right. Then the grid Γ contains ρ pairwise disjoint paths P_1, \dots, P_ρ such that, for every $h \in [\rho]$, the endpoints of P_h are p_h^{up} and p_h^{down} .*

Proof. We use induction on ρ . Clearly the lemma is obvious when $\rho = 0$. Let $(i, j) \in [k]^2$ such that p_ρ^{up} (resp. p_ρ^{down}) is the i -th (resp. j -th) vertex of the higher (lower) horizontal line counting from left to right. We examine first the case where $i \geq j$. Let P_ρ be the path created by starting from p_ρ^{up} , moving $k' - 1$ edges down, 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 p_h^{down} and moving one edge up (clearly, $P_h^{(\text{down})'}$ consists of a single edge). We also denote by $p_i^{(\text{down})'}$ the other endpoint of $P_i^{(\text{down})'}$. 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. By construction, none of the edges or vertices of P_ρ belongs in 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. $\{p_1^{(\text{down})'}, \dots, p_{\rho-1}^{(\text{down})'}\}$). From the induction hypothesis, Γ' contains $\rho - 1$ pairwise disjoint paths $P'_1, \dots, P'_{\rho-1}$ such that for every $h \in [\rho - 1]$, the endpoints of P_h are p_h^{up} and $p_h^{(\text{down})'}$. 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 $i < j$, just reverse the same grid upside down and the proof is identical (see Figure 11). \square

Lemma 8. *Let Γ be a $(k \times k)$ -grid embedded in the plane and assume that the vertices of its outer cycle, arranged in clockwise order, are:*

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

Let also H be a graph whose vertices have degree 0 or 1 and they 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. Let V^1 be the vertices of H that have degree 1 and let $H^1 = H[V^1]$. Then H^1 is a topological minor of Γ 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 $D = \{x_1^{\text{down}}, \dots, x_k^{\text{down}}\} \cup V^1$. We define ϕ_0 as in the statement of the lemma. 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 *crossing edges* E_C that have one endpoint in U and one in D . As $|V(H^1)| \leq 2k$ we obtain 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 $\{x_i^{\text{up}}, x_j^{\text{up}}\}$ in E_U as follows: it is 0 if there is no edge of E_U with an endpoint in $\{x_{i+1}^{\text{up}}, \dots, x_{j-1}^{\text{up}}\}$ and is $i > 0$ if the maximum depth of an edge with an endpoint in $\{x_{i+1}^{\text{up}}, \dots, x_{j-1}^{\text{up}}\}$ is $i - 1$. The *depth* of an edge $\{x_i^{\text{down}}, x_j^{\text{down}}\}$ is defined analogously. It directly follows, by the definition of depth that:

$$q^{\text{up}} = \max\{\text{depth}(e) \mid e \in E_U\} + 1 \leq |E_U| \quad (3)$$

$$q^{\text{down}} = \max\{\text{depth}(e) \mid e \in E_D\} + 1 \leq |E_D| \quad (4)$$

We now continue with the definition of ϕ_1 as follows:




Figure 12: An example of the proof of Lemma 8. On the left, the (16×16) -grid G is depicted along with the way the graph H (depicted on the right) is (partially) routed in it. In the figure $q^{\text{up}} = 3$, $q^{\text{down}} = 2$, $k' = 11$, and $\rho = 5$.

- for every edge $e = \{x_i^{\text{up}}, x_j^{\text{up}}\}$ in E_U , of depth l and such that $i < j$, let $\phi_1(e)$ be the path defined if we start in the grid G from v_i^{up} , move l steps down, then $j - i$ steps to the right, and finally move l steps up to the vertex v_j^{up} (by “number of steps” we mean number of edges traversed).
- for every edge $e = \{x_i^{\text{down}}, x_j^{\text{down}}\}$ in E_D , of depth l and such that $i < j$, let $\phi_1(e)$ be the path defined if we start in the grid G from v_i^{down} , move l steps up, then $j - i$ steps to the right, and finally move l steps down to the vertex v_j^{down} .

Notice that the above two steps define the values of ϕ_1 for all the upper and down edges. The construction guarantees that all paths in $\phi_1(E_U \cup E_D)$ are mutually non-crossing. Also, the distance between $\phi_0(U)$ and some horizontal line of Γ that contains edges of the images of the upper edges is $\max\{\text{depth}(e) \mid e \in E_U\}$ that, from (3), is equal to $q^{\text{up}} - 1$. Symmetrically, using (4) instead of (3), 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 $q^{\text{down}} - 1$. As a consequence, the graph

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

is a $(k \times k')$ -grid Γ' , where $k' = k - (q^{\text{up}} + q^{\text{down}})$, 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 q^{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 q^{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 (these endpoints are depicted as white vertices in the example of Figure 12). We use the notation $\{p_1^{\text{up}}, \dots, p_\rho^{\text{up}}\}$ (resp. $\{p_1^{\text{down}}, \dots, p_\rho^{\text{down}}\}$) 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 $\{p_1^{\text{up}}, \dots, p_\rho^{\text{up}}\}$ (resp. $\{p_1^{\text{down}}, \dots, p_\rho^{\text{down}}\}$) are ordered as they appear from left to right in the upper (lower) horizontal line of Γ' (this is possible because of the outerplanarity of H^+).

Notice that $\rho = |E(H^1)| - (|E_U| + |E_D|) \leq k - (|E_U| + |E_D|)$ which by (3) and (4) implies that $\rho \leq k'$.

As $\rho \leq k' \leq k$, we can now apply Lemma 7 on Γ' , $\{p_1^{\text{up}}, \dots, p_\rho^{\text{up}}\}$ and $\{p_1^{\text{down}}, \dots, p_\rho^{\text{down}}\}$ and obtain a collection $\{P_e \mid e \in E_C\}$ of ρ pairwise disjoint paths in Γ' between the vertices of $\{p_e^{\text{up}} \mid e \in E_C\}$ and the vertices of $\{p_e^{\text{down}} \mid e \in E_C\}$. 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 (ϕ_1, ϕ_2) provides the claimed topological isomorphism. \square

Lemma 9. *Let G be a graph with a linkage L consisting of k paths. Let also $\mathcal{U} = (\mathcal{X}, \mathcal{Z})$ be an L -tidy tilted grid of G with capacity m . Let also Δ be the closed-interior of the perimeter of $G_{\mathcal{U}}$. If $m > 2^k$, then G contains a linkage L' such that*

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

Proof. We use the notation $\mathcal{X} = \{X_1, \dots, X_m\}$ and $\mathcal{Z} = \{Z_1, \dots, Z_m\}$. Let $G_{\mathcal{U}}$ be the realization of \mathcal{U} in G and let G^* (resp. L^*) be the graph (resp. linkage) obtained from G

(resp. L) if we contract all edges in the paths of $\bigcup_{(i,j) \in \{1, \dots, r\}^2} I_{i,j}$, where $I_{i,j} = X_i \cap Z_j$, $i, j \in \{1, \dots, m\}$. We also define \mathcal{X}^* and \mathcal{Z}^* by applying the same contractions to their paths. Notice that $\mathcal{U}^* = (\mathcal{X}^*, \mathcal{Z}^*)$ is an L^* -tidy tilted grid of G^* with capacity m and that the lemma follows if we find a linkage L'^* such that the above three conditions are true for Δ^*, L^*, L'^* , and \mathcal{Z}^* , where Δ^* is the closed-interior of the perimeter of $G_{\mathcal{U}^*}^*$ (recall that $G_{\mathcal{U}^*}^*$ is the representation of \mathcal{U} that is isomorphic to $G_{\mathcal{U}^*}$).

Let $G^{*-} = (G^* \setminus \Delta^*) \cup \mathbf{UZ}$ and apply Proposition 3 on G^{*-} , Δ^* , and L^* . Let \mathcal{N} be a collection of strictly less than m 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 $L_0 = R \cup \mathcal{N}$ is a linkage of the graph $(G^* \setminus \Delta^*) \cup \mathbf{UN}$ 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 $(m \times m)$ -grid Γ is a topological minor of $G_{\mathcal{U}^*}^*$ via some pair (χ_0, χ_1) satisfying the conditions A and B in the definition of tilted grid.

We are now in position to apply Lemma 8 for the $(m \times m)$ -grid Γ and H . We obtain that $H^1 = L_0 \cap \Delta^*$ is a 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

$$Q = \bigcup_{e \in E(L)} \chi_1(e)$$

which, in turn, is a subgraph of $G_{\mathcal{U}^*}^*$. Observe that $L'^* = R \cup Q$ is a 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 $|\mathcal{N}| < m$, $E(\mathbf{UZ}^* \cap Q)$ is a proper subset of $E(\mathbf{UZ}^*)$. By construction of L'^* , it holds that $E(\mathbf{UZ} \cap L'^*) = E(\mathbf{UZ} \cap Q)$. Moreover, as $\mathcal{U}^* = (\mathcal{X}^*, \mathcal{Z}^*)$ is an L^* -tidy tilted grid of G^* , it follows that $E(\mathbf{UZ}^*) = E(\mathbf{UZ}^* \cap L^*)$. Therefore, Condition 3 follows. \square

3.7 Existence of an irrelevant vertex

We now bring together all results from the previous subsections in order to prove Theorem 1.

Lemma 10. *There exists an algorithm that, given an instance $(G, \mathcal{P} = \{(s_i, t_i) \in V(G)^2, i \in \{1, \dots, k\}\})$ of PDPP, either outputs a tree-decomposition of G of width at most $9 \cdot (k \cdot 2^{k+2} + 1) \cdot \lceil \sqrt{2k+1} \rceil$ or outputs an irrelevant vertex $x \in V(G)$ for (G, \mathcal{P}) . This algorithm runs in $2^{2^{O(k)}} \cdot n$ steps.*

Proof. Let $T = \{s_1, \dots, s_k, t_1, \dots, t_k\}$. By applying the algorithm of Lemma 1, for $r = k \cdot 2^{k+2}$ either we output a tree-decomposition of G of width at most $9(r+1) \cdot \lceil \sqrt{2k+1} \rceil$

or we find an internally chordless cycle C of G such that G contains a tight sequence of cycles $\mathcal{C} = \{C_0, \dots, C_r\}$ in G where $C_0 = C$ and all vertices of T are in the open exterior of C_r . From Lemma 1, this can be done in $2^{(r \cdot \sqrt{|T|})^{O(1)}} \cdot n = 2^{2^{O(k)}} \cdot n$ steps.

Assume that G has a linkage whose pattern is \mathcal{P} and, among all such linkages, let L be a \mathcal{C} -cheap one. Our aim is to prove that $V(L \cap C_0) = \emptyset$, i.e., we may pick x to be any of the vertices in D_0 .

First, we can assume that $k \geq 2$. Otherwise, if $k = 1$, the fact that L is \mathcal{C} -cheap, implies that $L \cap D_{r-1} = \emptyset \Rightarrow L \cap D_0 = \emptyset$ and we are done.

For every $i \in \{0, \dots, r\}$, we define $\mathcal{Q}^{(i)} = (\mathcal{C}^{(i)}, L^{(i)})$ where $\mathcal{C}^{(i)} = \{C_0, \dots, C_i\}$ and $L^{(i)}$ is the subgraph of L consisting of the union of the connected components of L that have common points with D_i . As $r + 1 > k$, at least one of $\mathcal{Q}^{(i)}, i \in \{0, \dots, r\}$ is touch-free. Let $\mathcal{Q}' = (\mathcal{C}', L')$ be the touch-free CL-configuration in $\{\mathcal{Q}^{(1)}, \dots, \mathcal{Q}^{(r)}\}$ of the highest index, say h . In other words, $\mathcal{C}' = \mathcal{C}^{(h)}$ and $L' = L^{(h)}$. Moreover, \mathcal{C}' is tight in G and L' is \mathcal{C}' -cheap. Let k' be the number of connected components of L' . We set $d = r - h$ and observe that $k' \leq k - d$, while \mathcal{C}' has $r' = r + 1 - d > 0$ concentric cycles. Again, we assume that $k' \geq 2$ as, otherwise, the fact that L' is \mathcal{C}' -cheap implies that $L' \cap D_{r'-1} = \emptyset \Rightarrow L' \cap D_0 = \emptyset$ and we are done. Therefore $0 \leq d \leq k - 2$.

As \mathcal{C}' is tight in G and L' is \mathcal{C}' -cheap, by Lemma 3, \mathcal{Q}' is convex. To prove that $V(L \cap C_0) = \emptyset$ it is enough to show that all segments of \mathcal{Q} have positive eccentricity and for this it is sufficient to prove that all segments of \mathcal{Q}' have positive eccentricity. Assume to the contrary that some segment P_0 of \mathcal{Q}' has eccentricity 0. Then, from the third condition in the definition of convexity we can derive the existence of a sequence $P_0, \dots, P_{r'-1}$ of segments such that for each $i \in \{0, \dots, r' - 1\}$, P_{i+1} is inside the zone of P_i . This implies the existence in the segment tree $T(\mathcal{Q}')$ of a path of length r' from its root to one of its leaves, therefore $T(\mathcal{Q}')$ has height r' . By Lemma 5, the real height of $T(\mathcal{Q}')$ is at most $2k' - 3$. By Observation 4, the dilation of $T(\mathcal{Q}')$ is at least $\frac{r'}{2k'-3} \geq \frac{k \cdot 2^{k+2} - d}{2k - 2d} > \frac{k \cdot 2^{k+2}}{2k} = 2^{k+1}$. By Observation 3 and Lemma 6, G contains an L' -tidy tilted grid $\mathcal{U} = (\mathcal{X}, \mathcal{Z})$ of capacity $> 2^k$. From Lemma 9, G contains another linkage L'' with the same pattern as L' and such that $c(L'') < c(L')$, a contradiction to the fact that L' is \mathcal{C}' -cheap.

Since $V(L \cap C_0) = \emptyset$, any vertex of G in the closed-interior of C_0 is irrelevant. \square

Proof of Theorem 1. The proof follows from Lemma 10, taking into account that, for every $k \geq 1$,

$$82 \cdot k^{3/2} \cdot 2^k > 9 \cdot (k \cdot 2^{k+2} + 1) \cdot \lceil \sqrt{2k + 1} \rceil$$

\square

4 An algorithm for PDPP

In this section we prove Theorem 2. In particular, we briefly describe an algorithm that, given an instance (G, \mathcal{P}) of DPP where G is planar, provides a solution to PDPP, if

one exists, in $2^{2^{O(k)}} \cdot n^{O(1)}$ steps.

Our algorithm is based on the following proposition.

Proposition 4 ([25]). *There exists an algorithm that, given an instance (G, \mathcal{P}) of PDPP and a tree decomposition of G of width at most w , either reports that (G, \mathcal{P}) is a NO-instance or outputs a solution of PDPP for (G, \mathcal{P}) in $2^{O(w \log w)} \cdot n$ steps.*

Proof of Theorem 2. By applying the algorithm of Lemma 10, we either find an irrelevant vertex v for (G, \mathcal{P}) or we obtain a tree-decomposition of G of width $2^{O(k)}$. In the first case, we again look for an irrelevant vertex in the equivalent instance $(G, \mathcal{P}) \leftarrow (G \setminus v, \mathcal{P})$. This loop breaks if the second case appears, namely when a tree decomposition of G of width $2^{O(k)}$ is found. Then we apply the algorithm of Proposition 4, that solves the problem in $2^{2^{O(k)}} \cdot n$ steps. As, unavoidably, the loop will break in less than n steps, the claimed running time follows. \square

Acknowledgment. We thank Ken-ichi Kawarabayashi and Paul Wollan for providing details on the bounds in [14]. We are particularly thankful to the two anonymous reviewers for their detailed and insightful reviews.

References

- [1] Isolde Adler, Frederic Dorn, Fedor V. Fomin, Ignasi Sau, and Dimtrios M. Thilikos. Faster parameterized algorithms for minor containment. In *12th Scandinavian Workshop on Algorithm Theory—SWAT 2010 (Bergen)*, volume 6139 of *Lecture Notes in Computer Science*, pages 322 – 333. Springer, Berlin, 2010.
- [2] Isolde Adler, Stavros G. Kolliopoulos, Philipp Klaus Krause, Daniel Lokshtanov, Saket Saurabh, and Dimitrios M. Thilikos. Tight bounds for linkages in planar graphs. In *Automata, Languages and Programming - 38th International Colloquium, ICALP 2011*, volume 6755 of *Lecture Notes in Computer Science*, pages 110–121. Springer, Berlin, 2011.
- [3] Isolde Adler, Stavros G. Kolliopoulos, and Dimitrios M. Thilikos. Planar disjoint-paths completion. *Algorithmica*, pages 1–25, 2015.
- [4] Hans L. Bodlaender. A linear time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on Computing*, 25:1305–1317, 1996.
- [5] Marek Cygan, Dániel Marx, Marcin Pilipczuk, and Michal Pilipczuk. The planar directed k -vertex-disjoint paths problem is fixed-parameter tractable. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 197–206, 2013.

- [6] Anuj Dawar, Martin Grohe, and Stephan Kreutzer. Locally excluding a minor. In *LICS'07*, pages 270–279. IEEE Computer Society, 2007.
- [7] Anuj Dawar and Stephan Kreutzer. Domination problems in nowhere-dense classes. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2009, December 15-17, 2009, IIT Kanpur, India*, pages 157–168, 2009.
- [8] Jeff Erickson and Amir Nayyeri. Shortest non-crossing walks in the plane. In *22nd Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, 2011*, pages 297–208. ACM-SIAM, 2011.
- [9] Petr A. Golovach, M. Kamiński, D. Paulusma, and D. M. Thilikos. Induced packing of odd cycles in a planar graph. In *20th International Symposium on Algorithms and Computation, ISAAC 2009*, volume 5878 of *LNCS*, pages 514–523. Springer, Berlin, 2009.
- [10] Qian-Ping Gu and Hisao Tamaki. Improved bounds on the planar branchwidth with respect to the largest grid minor size. *Algorithmica*, 64(3):416–453, 2012.
- [11] Illya V. Hicks. Branch decompositions and minor containment. *Networks*, 43(1):1–9, 2004.
- [12] Ken-ichi Kawarabayashi and Yusuke Kobayashi. The induced disjoint path problem. In *13th Conference on Integer Programming and Combinatorial Optimization, IPCO 2008*, volume 5035 of *Lecture Notes in Computer Science*, pages 47–61. Springer, Berlin, 2008.
- [13] Ken-ichi Kawarabayashi and Bruce Reed. Odd cycle packing. In *42nd ACM Symposium on Theory of Computing, STOC 2010*, pages 695–704. ACM, 2010.
- [14] Ken-ichi Kawarabayashi and Paul Wollan. A shorter proof of the graph minor algorithm: the unique linkage theorem. In *42nd ACM Symposium on Theory of Computing, STOC 2010*, pages 687–694. ACM, 2010.
- [15] Yusuke Kobayashi and Ken-ichi Kawarabayashi. Algorithms for finding an induced cycle in planar graphs and bounded genus graphs. In *20th ACM-SIAM Symposium on Discrete Algorithms, SODA 2009*, pages 1146–1155. ACM-SIAM, 2009.
- [16] Mark R. Kramer and Jan van Leeuwen. The complexity of wire-routing and finding minimum area layouts for arbitrary VLSI circuits. *Advances in Comp. Research*, 2:129–146, 1984.
- [17] James F. Lynch. The equivalence of theorem proving and the interconnection problem. *ACM SIGDA Newsletter*, 5:31–36, 1975.

- [18] Matthias Middendorf and Frank Pfeiffer. On the complexity of the disjoint paths problem. *Combinatorica*, 13(1):97–107, 1993.
- [19] Ljubomir Perković and Bruce Reed. An improved algorithm for finding tree decompositions of small width. In Peter Widmayer, Gabriele Neyer, and Stephan Eidenbenz, editors, *Graph-Theoretic Concepts in Computer Science*, volume 1665 of *Lecture Notes in Computer Science*, pages 148–154. Springer, Berlin, 1999.
- [20] Bruce A. Reed, Neil Robertson, Alexander Schrijver, and Paul D. Seymour. Finding disjoint trees in planar graphs in linear time. In *Graph Structure Theory, Proceedings of a AMS-IMS-SIAM Joint Summer Research Conference on Graph Minors held June 22 to July 5, 1991, at the University of Washington, Seattle*, pages 295–301. AMS-IMS-SIAM, 1991.
- [21] Neil Robertson and Paul D. Seymour. Graph Minors. X. Obstructions to Tree-decomposition. *J. Comb. Theory Ser. B*, 52(2):153–190, 1991.
- [22] Neil Robertson and Paul D. Seymour. Graph Minors. XIII. The disjoint paths problem. *J. Comb. Theory Ser. B*, 63(1):65–110, 1995.
- [23] Neil Robertson and Paul D. Seymour. Graph minors. XXI. Graphs with unique linkages. *J. Comb. Theory Ser. B*, 99(3):583–616, 2009.
- [24] Neil Robertson and Paul D. Seymour. Graph minors. XXII. Irrelevant vertices in linkage problems. *J. Comb. Theory, Ser. B*, 102(2):530–563, 2012.
- [25] Petra Scheffler. A practical linear time algorithm for disjoint paths in graphs with bounded tree-width. Technical Report 396/1994, FU Berlin, Fachbereich 3 Mathematik, 1994.
- [26] Alexander Schrijver. Finding k disjoint paths in a directed planar graph. *SIAM J. Comput.*, 23(4):780–788, 1994.
- [27] Alexander Schrijver. *Combinatorial optimization. Polyhedra and efficiency. Vol. A*. Springer-Verlag, Berlin, 2003.
- [28] Jens Vygen. NP-completeness of some edge-disjoint paths problems. *Discrete Appl. Math.*, 61(1):83–90, 1995.