

Bounds on Half Graph Orders in Powers of Sparse Graphs*

Marek Sokołowski[†]

Institute of Informatics, University of Warsaw

marek.sokolowski@mimuw.edu.pl

Submitted: Feb 21, 2022; Accepted: Dec 23, 2022; Published: Apr 7, 2023

© The author. Released under the CC BY-ND license (International 4.0).

Abstract

Half graphs and their variants, such as semi-ladders and co-matchings, are configurations that encode total orders in graphs. Works by Adler and Adler (Eur. J. Comb.; 2014) and Fabiański et al. (STACS; 2019) prove that in powers of sparse graphs, one cannot find arbitrarily large configurations of this kind. However, these proofs either are non-constructive, or provide only loose upper bounds on the orders of half graphs and semi-ladders. In this work we provide nearly tight asymptotic lower and upper bounds on the maximum order of half graphs, parameterized by the power, in the following classes of sparse graphs: planar graphs, graphs with bounded maximum degree, graphs with bounded pathwidth or treewidth, and graphs excluding a fixed clique as a minor.

The most significant part of our work is the upper bound for planar graphs. Here, we employ techniques of structural graph theory to analyze semi-ladders in planar graphs via the notion of cages, which expose a topological structure in semi-ladders. As an essential building block of this proof, we also state and prove a new structural result, yielding a fully polynomial bound on the neighborhood complexity in the class of planar graphs.

Mathematics Subject Classifications: 05C12, 05C35, 05C75

1 Introduction

It is widely known that there is a huge array of algorithmic problems deemed to be computationally hard. One of the ways of circumventing this issue is limiting the set of possible instances of a problem by assuming a more manageable structure. For example, restricting our attention to graph problems, we can exploit the planarity of the graph

*The results of this paper have been presented in the master thesis, defended by the author at the University of Warsaw. The extended abstract of this work has been presented at EUROCOMB 2021 [39].

[†]This work is a part of project TOTAL that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 677651).

instances through the planar separator theorem [19] or Baker’s technique [3]. Analogously, if graphs have bounded treewidth or pathwidth, we can solve many hard problems by means of dynamic programming on tree or path decompositions [4]. These examples present some of the algorithmic techniques which allow us to utilize the structural sparsity of graph instances.

Nešetřil and Ossona de Mendez have proposed two abstract notions of sparsity in graphs: bounded expansion [23] and nowhere denseness [27, 28]. Intuitively speaking, a class \mathcal{C} of graphs has bounded expansion if one cannot obtain arbitrarily dense graphs by picking a graph $G \in \mathcal{C}$ and contracting pairwise disjoint connected subgraphs of G with fixed radius to single vertices. More generally, \mathcal{C} is nowhere dense if one cannot produce arbitrarily large cliques as graphs as a result of the same process; see Preliminaries (Section 2) for exact definitions. These combinatorial notions turn out to be practical in the algorithm design. In fact, they are some of the fundamental concepts of Sparsity — a research area concerning the combinatorial properties and the algorithmic applications of sparse graphs. The tools of Sparsity allow us to design efficient algorithms in nowhere dense classes of graphs and in classes of graphs with bounded expansion for problems such as subgraph isomorphism [24, 28] or distance- d dominating set [5, 7, 10, 17]. In fact, under reasonable complexity assumptions, nowhere denseness exactly characterizes the classes of graphs closed under taking subgraphs that allow efficient algorithms for all problems definable in first-order logic [9, 15]. For a comprehensive introduction to Sparsity, we refer to the book by Nešetřil and Ossona de Mendez [29] and to the lecture notes from the University of Warsaw [33].

Research in Sparsity provided a plethora of technical tools for structural analysis of sparse graphs, many of which are in the form of various graph parameters. For example, the classes of graphs with bounded expansion equivalently have bounded **generalized coloring numbers** — the **weak** and **strong d -coloring numbers** are uniformly bounded in a class \mathcal{C} of graphs for each fixed $d \in \mathbb{N}$, if and only if \mathcal{C} has bounded expansion [44]. Another example is the concept of **p -centered colorings** [25], which have been shown to require a bounded number of colors for each fixed $p \in \mathbb{N}$ exactly for the classes of graphs with bounded expansion [26]. It is interesting and useful to determine the asymptotic growth (depending on the parameter; d or p in the examples above) of these measures in well-studied classes of sparse graphs, such as planar graphs, graphs with bounded pathwidth or treewidth, proper minor-closed classes of graphs etc., as it should lead to a better understanding of these classes. For instance, the asymptotic behavior of generalized coloring numbers has been studied in proper minor-closed classes of graphs and in planar graphs [14, 22, 42]. Also, there has been a recent progress on p -centered colorings in proper minor-closed classes, in planar graphs, and in graphs with bounded maximum degree [6, 8, 34].

In this work, we will consider another kind of structural measures that behave nicely in sparse graphs, which are related to the concepts of half graphs, semi-ladders, and co-matchings.

Definition 1. In an undirected graph $G = (V, E)$, for an integer $\ell \geq 1$, 2ℓ different vertices $a_1, a_2, \dots, a_\ell, b_1, b_2, \dots, b_\ell$ form:

- a **half graph** of order ℓ if for each pair of indices $i, j \in [1, \ell]$, we have $\{b_i, a_j\} \in E$ if and only if $i < j$ (Figure 1(a));

- a **semi-ladder** of order ℓ if we have $\{b_i, a_j\} \in E$ for each pair of indices i, j such that $1 \leq i < j \leq \ell$, and $\{b_i, a_i\} \notin E$ for each $i \in [1, \ell]$ (Figure 1(b));
- a **co-matching** of order ℓ if for each pair of indices $i, j \in [1, \ell]$, we have $\{b_i, a_j\} \in E$ if and only if $i \neq j$ (Figure 1(c)).

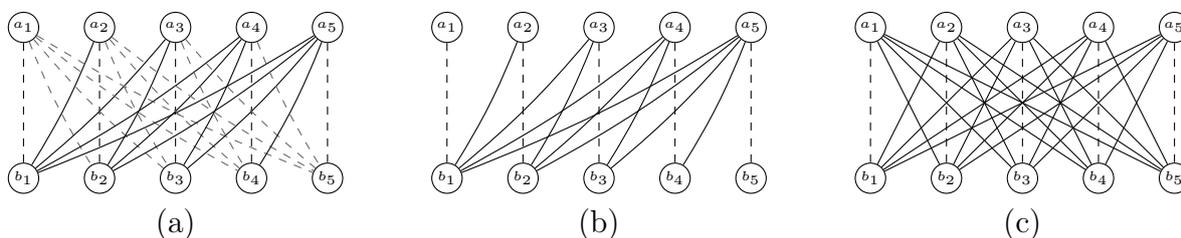


Figure 1: The configurations in Definition 1. Solid lines indicate pairs of vertices connected by an edge, and dashed edges indicate pairs of vertices not connected by an edge. Intuitively, half graphs and co-matchings encode the $<$, and \neq relations, respectively, on pairs $(a_1, b_1), (a_2, b_2), \dots, (a_\ell, b_\ell)$ of vertices.

Naturally, each half graph, and co-matching is also a semi-ladder. Hence, if for a class \mathcal{C} of graphs, the orders of semi-ladders occurring in any graph in \mathcal{C} are bounded from above by some constant M , then M is also the corresponding upper bound for orders of half graphs and co-matchings. We also remark that there is a simple proof utilizing Ramsey’s theorem [12, Lemma 2] which demonstrates that if an arbitrary class \mathcal{C} of graphs has a uniform upper bound on the orders of half graphs in any graph in \mathcal{C} , and a separate uniform upper bound on the orders of co-matchings, then there also exists a uniform upper bound on the maximum order of semi-ladders in \mathcal{C} .

The absence of arbitrarily large half graphs in a given class of graphs is a good indicator of the well-structuredness of the class, to the point where it forms a basis of a key definition in model theory. Namely, we say that a class \mathcal{C} of graphs (or relational structures) is **stable** if every first-order interpretation of \mathcal{C} has a uniform bound on the maximum order of half graphs. Formally, for a fixed first-order formula $\varphi(\bar{x}, \bar{y})$, where \bar{x} and \bar{y} are tuples of variables, and a graph G , we define the bipartite graph $G^\varphi = (V_1, V_2, E)$ as follows:

$$\begin{aligned} V_1(G^\varphi) &= V(G)^{\bar{x}}, \\ V_2(G^\varphi) &= V(G)^{\bar{y}}, \\ E(G^\varphi) &= \{\{\bar{a}, \bar{b}\} \mid G \models \varphi(\bar{a}, \bar{b}), \bar{a} \in V_1(G^\varphi), \bar{b} \in V_2(G^\varphi)\}. \end{aligned}$$

Then, \mathcal{C} is stable if and only if for each formula φ , the class $\mathcal{C}^\varphi = \{G^\varphi \mid G \in \mathcal{C}\}$ admits a uniform bound on the maximum order of half graphs.

Stability is one of important properties of structures in mathematical logic, and has been studied extensively — see the book by Pillay [36] for the general exposition of the topic, and the work of Pilipczuk et al. [35] for the information on connections of the stability theory with sparse graphs. As an example use of stability in graph theory, one can consider stable classes of graphs with some additional structural properties. For instance, stable classes of graphs of bounded cliquewidth, linear cliquewidth or twin-width

can, in fact, be constructed from the classes of graphs with bounded treewidth, pathwidth and sparse twin-width, respectively, by means of some first-order transformations [13, 30, 31]. Also, in the classes of graphs with a uniform bound on the orders of half graphs, Malliaris and Shelah improved the bounds on the parameters in the statement of the Szemerédi’s regularity lemma [20].

Notably, every nowhere dense class of graphs \mathcal{C} is stable. This has been proved by Adler and Adler [1], who applied the ideas of Podewski and Ziegler [37]. While the proof of Adler and Adler relies on the compactness theorem of first-order logic, Pilipczuk, Siebertz, and Toruńczyk proposed a direct proof of this fact, giving explicit upper bounds on the orders of half graphs [35].

In this work, we will focus on specific interpretations of nowhere dense classes of graphs — namely, the powers of graphs. For an undirected graph G and an integer d , we define the graph G^d as follows:

$$V(G^d) = V(G) \quad \text{and} \quad E(G^d) = \{\{u, v\} \mid \text{dist}_G(u, v) \leq d\}.$$

Then, the d -th power of a class \mathcal{C} is defined as $\mathcal{C}^d = \{G^d \mid G \in \mathcal{C}\}$.

Note that it is immediate that for any nowhere dense classes \mathcal{C} of graphs, the orders of half graphs in graphs from \mathcal{C}^d are uniformly bounded for each $d \in \mathbb{N}$: it simply follows from the stability of \mathcal{C} . However, in this specific case, it turns out that the same property holds for all the configurations from Definition 1:

Theorem 2 ([1, 12]). *For every nowhere dense class \mathcal{C} of graphs, there exists a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that for every positive integer d , the orders of semi-ladders in graphs from \mathcal{C}^d are uniformly bounded from above by $f(d)$.*

We remark that Theorem 2 directly bounds the orders of half graphs and co-matchings found in graphs from \mathcal{C}^d as well.

An important motivation for studying semi-ladders is a work of Fabiański et al. [12], who presented a fixed-parameter algorithm solving the DISTANCE- d DOMINATING SET problem in sparse graphs:

DISTANCE- d DOMINATING SET

FIXED: a nowhere dense class \mathcal{C} of graphs, distance $d \geq 1$

INPUT: a graph $G \in \mathcal{C}$, a parameter $k \geq 1$

OUTPUT: a subset $D \subseteq V(G)$, $|D| \leq k$, such that each vertex of G is at distance at most d from a vertex of D , provided such a set exists.

In any nowhere dense class \mathcal{C} , the running time of this algorithm is bounded from above by $2^{O(k \log k)} \cdot (n + m)$ for each fixed $d \in \mathbb{N}$. If \mathcal{C} has bounded expansion, we can infer a more concrete upper bound: $O((\xi \cdot k^L)^k + k^L \cdot (n + m))$, where L is the bound on the order of semi-ladders in graphs from \mathcal{C}^d (i.e., the bound $f(d)$ from Theorem 2), and ξ is a constant related to the concept of neighborhood complexity of \mathcal{C} :

Theorem 3 ([38]). *Let \mathcal{C} be a class of graphs with bounded expansion, and let $d \geq 1$ be an integer. For any graph G and its vertex v , let $N^d[v]$ be the closed d -neighborhood of v ; that is, the set of vertices at distance at most d from v . Then, for every non-empty set X*

of vertices of G , the number of different intersections of X with a closed d -neighborhood of any vertex of G is linear in the size of X :

$$|\{N^d[v] \cap X \mid v \in V(G)\}| \leq \xi \cdot |X|, \quad (1)$$

where ξ is a constant depending only on \mathcal{C} and d . The left-hand side of (1) is called the *neighborhood complexity* of G .

If \mathcal{C} is nowhere dense, the runtime guarantees of this algorithm are slightly looser, but they still rely heavily on the value of L . Hence, if we find tight asymptotic bounds on L as a function of d in the d -th powers of various concrete classes of sparse graphs, we should be able to infer concrete bounds on the running time of the algorithm of Fabiański et al. in these classes.

As mentioned above, the value of L is bounded in powers of nowhere dense classes of graphs. However, the proof in [12] relies on the description of nowhere denseness through uniform quasi-wideness (see Definition 112). Hence, the resulting bounds on the maximum order of semi-ladders are usually suboptimal — for instance, applying them directly to the d -th power of the class of planar graphs yields a $d^{O(d^6)}$ upper bound; compare this to a $d^{O(d)}$ bound we present in this work (Theorem 52). Therefore, it is natural to pose a question: can we find good estimates on the maximum orders of half graphs and semi-ladders occurring in the powers of well-studied classes of sparse graphs?

Our results. We state and prove asymptotic lower and upper bounds on the orders of semi-ladders in the d -th powers of the following classes of sparse graphs: graphs with maximum degree bounded by Δ , planar graphs, graphs with pathwidth bounded by p , graphs with treewidth bounded by t , and graphs excluding the complete graph K_t as a minor (Figure 2). We reiterate that each bound on the orders of semi-ladders directly translates to a corresponding bound on the orders of half graphs and co-matchings. Our results are asymptotically almost tight — they allow us to understand how quickly the maximum orders of half graphs grow, as functions of d , in these classes of graphs.

In Section 3, we describe three constructions proving the lower bounds stated in Figure 2. Each of the constructions, for an integer d , creates a graph whose d -th power contains a large half graph. (Again, these graphs also contain large semi-ladders.) Conversely, the remaining part of this work (Sections 4, 5, 6) contains proofs of the upper bounds on the maximum orders of semi-ladders in the considered classes of graphs; these proofs also yield the upper bounds on the orders of half graphs and co-matchings.

In Section 4, we prove that the neighborhood complexity in planar graphs has an upper bound which is polynomial both in $|X|$ and in d . We note that Theorem 3 already proves a bound linear in $|X|$; however, the proof of Reidl et al. [38] requires an exponential dependence on d , even in the class of planar graphs. The polynomial dependence on d asserted by our result is crucial in the proof of the upper bound on semi-ladders in the class of planar graphs.

Section 5 is fully devoted to the proof of the upper bound on the maximum order of a semi-ladder in the d -th power of any planar graph. The proof is quite involved and should be considered the main contribution of this work. This proof works as follows: we start with a huge semi-ladder in the d -th power of a chosen planar graph G . Using this semi-ladder, we find in G configurations with more and more structure, which we call: quasi-cages, cages, ordered cages, identity ordered cages, neighbor cages, and separating cages,

Class of graphs	Lower bound	Upper bound
DEGREE $\leq \Delta$	$\Delta^{\Omega(d)}$ (Cor. 7)	$\Delta^d + 1$ (Thm. 99)
PLANAR	$2^{\lceil \frac{d}{2} \rceil}$ (Cor. 9)	$d^{O(d)}$ (Thm. 52)
PATHWIDTH $\leq p$	$d^{p-O(1)}$ (Cor. 11)	$(dp)^{O(p)}$ (Cor. 110)
TREEWIDTH $\leq t$	$2^{d^{\Omega(t)}}$ (Cor. 23)	$d^{O(d^{t+1})}$ (Cor. 119)
K_t -MINOR-FREE	$2^{d^{\Omega(t)}}$ (Cor. 25)	$d^{O(d^{t-1})}$ (Cor. 118)

Figure 2: Bounds on the maximum half graphs and semi-ladders in the d -th powers of classes of graphs proved in this work.

in this order. Each extraction step requires us to forfeit a fraction of the configuration, but we give good estimates on the maximum size loss incurred in the process. Eventually, we derive a direct upper bound on the maximum order of a separating cage that can be found in a planar graph. Retracing the extraction steps, we find a uniform upper bound on the maximum order of any semi-ladder in the d -th power of a planar graph.

Section 6 presents the proofs of upper bounds in the powers of other considered classes of graphs: graphs with bounded degree, graphs with bounded pathwidth or treewidth, and K_t -minor-free graphs. In the case of graphs of bounded pathwidth, we generalize the sunflower lemma (Lemma 107), which ultimately helps us to extract structural properties from path decompositions of bounded width. In the proof of the upper bound in the class of graphs excluding K_t as a minor, we rely on weak coloring numbers, and upper bounds on them in K_t -minor-free graphs proved by van den Heuvel et al. [42].

2 Preliminaries

Distances. In this work — unless specified — we only consider undirected graphs, without loops or multiple edges with the same endpoints. In a graph G , for two of its vertices u, v (not necessarily different), we denote the distance between these two vertices in G as $\text{dist}_G(u, v)$; if these vertices are not in the same connected component of G , we set $\text{dist}_G(u, v) = +\infty$.

When the graph is clear from the context, we may omit the name of the graph from our notation and simply write $\text{dist}(u, v)$.

Different definitions of half graphs. In some applications (e.g. in the definition of the stability in the Introduction), half graphs are allowed to contain multiple occurrences of any single vertex. However, a straightforward greedy procedure can convert this variant

of half graphs into the half graphs from Definition 1 of order at least half of the original order. Since we will only consider an asymptotic behavior of the orders of half graphs and other structures in various classes of graphs, we can safely assume the variant of half graphs from Definition 1.

Distance- d half graphs (semi-ladders, co-matchings). In many parts of this work, it will be more useful to operate on the original graph G instead of its d -th power G^d — some of the objects and decompositions defined later are properly defined in the original classes of graphs, but not in their powers. Hence, we define **distance- d** half-graphs, semi-ladders, and co-matchings that exist in G if and only if their counterparts from Definition 1 exist in G^d .

In an undirected graph $G = (V, E)$, for a pair of integers $d, \ell \geq 1$, we say that 2ℓ different vertices $a_1, a_2, \dots, a_\ell, b_1, b_2, \dots, b_\ell$ form:

- a **distance- d half graph** of order ℓ if for each pair of indices i, j such that $i, j \in [1, \ell]$, we have $\text{dist}(b_i, a_j) \leq d$ if and only if $i \leq j$;
- a **distance- d semi-ladder** of order ℓ if we have $\text{dist}(b_i, a_j) \leq d$ for each pair of indices i, j such that $1 \leq i < j \leq \ell$, and $\text{dist}(b_i, a_j) > d$ for each $i \in [1, \ell]$;
- a **distance- d co-matching** of order ℓ if for each pair of indices i, j such that $i, j \in [1, \ell]$, we have $\text{dist}(b_i, a_j) \leq d$ if and only if $i \neq j$.

Pathwidth, path decompositions. We define a path decomposition (W_1, W_2, \dots, W_k) of a graph G as any sequence of subsets of $V(G)$ (called **bags**) with the following properties:

- For each edge $uv \in E(G)$, there exists a bag W_i such that $u, v \in W_i$.
- For each vertex $v \in V(G)$, v appears in a contiguous non-empty subset of the sequence.

The width of the decomposition is $\max_{i=1}^k |W_i| - 1$. The **pathwidth** of a graph G , denoted $\text{pw}(G)$, is the minimum possible width of any path decomposition of G .

Treewidth, tree decompositions. A tree decomposition of a graph G is an undirected tree \mathcal{T} whose each vertex t is associated with a subset of vertices X_t , called the **bag** of t . The tree has the following properties:

- For each edge $uv \in E(G)$, there exists a vertex $t \in V(\mathcal{T})$ such that $u, v \in X_t$.
- For each vertex $v \in V(G)$, the set of vertices $\{t \mid v \in X_t\}$ forms a connected non-empty subtree of \mathcal{T} .

Similarly to the case of path decomposition, the width of a tree decomposition is defined as $\max_{t \in V(\mathcal{T})} |X_t| - 1$. The **treewidth** of a graph G , denoted $\text{tw}(G)$, is the minimum possible width of any tree decomposition of G .

Graph minors, minor-free classes of graphs. A graph H is a **minor** of G if one can map each vertex $v \in V(H)$ to a non-empty connected subgraph $A(v)$ of G so that every

two subgraphs $A(u), A(v)$ are vertex-disjoint, and for each edge $uv \in E(H)$, subgraphs $A(u)$ and $A(v)$ are connected by an edge. A graph H is a **depth- d minor** of G if additionally each subgraph $A(v)$ has radius not exceeding d .

A class \mathcal{C} of graphs is H -minor-free if no graph $G \in \mathcal{C}$ contains H as a minor. We note that for any class \mathcal{C} and graph H , if \mathcal{C} is $K_{|H|}$ -minor-free, then it is also H -minor-free. Hence, the maximum order of any semi-ladder in an H -minor-free class of graphs is bounded from above by the upper bound for $K_{|H|}$ -minor-free graphs.

As treewidth is a minor-closed parameter, and $\text{tw}(K_{t+2}) = t + 1$, the following holds:

Theorem 4. *Every graph G such that $\text{tw}(G) \leq t$ is also K_{t+2} -minor-free.*

Bounded expansion. We follow the notation used in the lecture notes from the University of Warsaw [33].

For a graph G and an integer $d \in \mathbb{N}$, we define $\nabla_d(G)$ to be the maximum edge density of any depth- d minor of G :

$$\nabla_d(G) := \sup \left\{ \frac{|E(H)|}{|V(H)|} \mid H \text{ is a depth-}d \text{ minor of } G \right\}.$$

A class \mathcal{C} of graphs has **bounded expansion** if there exists a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that we have $\nabla_d(G) \leq f(d)$ for each $G \in \mathcal{C}$, $d \in \mathbb{N}$.

Nowhere denseness. For a graph G and an integer $d \in \mathbb{N}$, we define $\omega_d(G)$ to be the maximum size of a clique which is a depth- d minor of G :

$$\omega_d(G) := \sup \{t \mid K_t \text{ is a depth-}d \text{ minor of } G\}.$$

A class \mathcal{C} of graphs is **nowhere dense** if there exists a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that we have $\omega_d(G) \leq f(d)$ for each $G \in \mathcal{C}$, $d \in \mathbb{N}$.

It can be easily verified that each class \mathcal{C} with bounded expansion is also nowhere dense. However, the converse statement is not true: there are nowhere dense classes which have unbounded expansion.

Paths. In an undirected graph G , a path is a sequence of vertices $P = (v_0, v_1, v_2, \dots, v_\ell)$ such that $(v_{i-1}, v_i) \in E(G)$ for each $i \in \{1, 2, \dots, \ell\}$. We say that P originates from v_0 and terminates at v_ℓ . The path is simple if all vertices $v_0, v_1, v_2, \dots, v_\ell$ are different.

For a simple path P and two vertices $u, v \in P$ with u occurring on the path before v , we define $P[u, v]$ as the subpath of P starting in u ending in v . Next, for any two paths $P = (p_0, p_1, \dots, p_x)$, $Q = (q_0, q_1, \dots, q_y)$ such that $p_x = q_0$, we denote by $P \cdot Q = (p_0, p_1, \dots, p_x, q_1, q_2, \dots, q_y)$ the concatenation of these two paths. Also, for any path $P = (v_0, v_1, \dots, v_\ell)$, we define $P^{-1} = (v_\ell, v_{\ell-1}, \dots, v_0)$ as the version of P with the order of vertices reversed, and $\text{len}(P) = \ell$ as its length.

Degenerate graphs. A k -degenerate graph ($k \geq 1$) [18] is a graph in which every subgraph contains a vertex of degree at most k . For instance, each planar graph is 5-degenerate.

In this paper, we use one property of degenerate graphs — every k -degenerate graph is $(k + 1)$ -colorable [41].

We remark that classes of degenerate graph are not necessarily nowhere dense (e.g. each subdivision of a graph is a 2-degenerate graph). Consequently, the class of 2-degenerate graphs contains arbitrary large semi-ladders, since subdivided half graphs are 2-degenerate.

Distance- d profiles. For a vertex $v \in V(G)$ and a subset of vertices $S \subseteq V(G)$, a **distance- d profile** of v on A in G is a function $\pi_d[v, S]$ from S to $\{0, 1, 2, \dots, d, +\infty\}$ that assigns each vertex from S either its distance from v , or $+\infty$ if this distance exceeds d .

3 Lower bounds on the orders of half graphs

In this section, we present a series of constructions which provide us with graphs belonging to the considered classes and containing large half graph. Since each half graph is also a semi-ladder, the constructed graphs will also contain large semi-ladders.

The section is divided into several subsections. In Subsection 3.1, we produce a family of graphs, each having bounded degree, but containing sizeable half graphs. It will later turn out that a subfamily of these graphs — specifically, those graphs in the family whose maximum degrees are bounded from above by 4 — consists solely of planar graphs. These graphs also witness an exponential lower bound on the maximum half graph size in the class of planar graphs.

In Subsection 3.2, we produce graphs with small pathwidth, yet large half graphs. Despite the construction being rather simple — a “core” containing $\Theta(p)$ vertices where p is the bound on pathwidth, together with a huge number of vertices connected only to the vertices of the core — it gives us a polynomial lower bound on the possible co-matching order, with the degree of the polynomial depending linearly on p .

In Subsection 3.3, we construct graphs with small treewidth and large half graphs. The construction is quite non-trivial — each graph in the family will be formed from multiple copies of smaller graphs in this family. This dependence allows us to create large half graphs in this class of graphs by means of induction. As the graphs with bounded treewidth do not contain large cliques as minors (Theorem 4), we will immediately find large half graphs in K_t -minor-free graphs.

3.1 Graphs with bounded degree and planar graphs

In this subsection, we will prove the following theorem:

Theorem 5. *For each $\Delta \geq 4$ and odd $d \geq 1$, there exists a graph with maximum degree bounded by Δ containing a distance- d half graph of order $\lfloor \frac{\Delta}{2} \rfloor^{\lceil \frac{d}{2} \rceil}$.*

We fix two integer variables, $k \geq 2$ and $h \geq 1$. We will construct an undirected graph $H_{k,h}$ with the following properties:

- The degree of each vertex in the graph is bounded by $2k$.
- The graph contains vertices $a_1, b_1, a_2, b_2, \dots, a_{k^h}, b_{k^h}$ forming a distance- $(2h-1)$ half graph of order k^h .

In the construction of $H_{k,h}$, we will create two isomorphic rooted trees \mathcal{A} and \mathcal{B} , each having k^h leaves, and connect pairs of vertices of these trees using paths of small, positive length.

We fix an alphabet $\Sigma = \{1, 2, \dots, k\}$ with the natural ordering of its elements. This allows us to introduce a lexicographical ordering of all words in Σ^* , denoted \preceq . For a word $s \in \Sigma^*$, by $|s|$ we mean the length of s . And for a pair of words $x, y \in \Sigma^*$, by $x \cdot y$ we mean the concatenation of x and y .

For each word $s \in \Sigma^*$, $|s| \in [0, h]$, we create two vertices: $a_s \in V(\mathcal{A})$ and $b_s \in V(\mathcal{B})$. The edges in both trees are constructed in the following way: for each word $s \in \Sigma^*$, $|s| \leq h - 1$, and each character $c \in \Sigma$, we add an edge between vertices a_s and $a_{s \cdot c}$ in \mathcal{A} , and between the vertices b_s and $b_{s \cdot c}$ in \mathcal{B} .

We also connect the trees in the following way: for each word $s \in \Sigma^*$, $|s| \leq h - 1$, and each pair of characters $c, d \in \Sigma$ such that $c < d$, we connect the vertices $a_{s \cdot d} \in V(\mathcal{A})$ and $b_{s \cdot c} \in V(\mathcal{B})$ by a path of length $2|s| + 1$.

This finishes the construction of $H_{k,h}$. Examples of such graphs can be found in Figure 3.

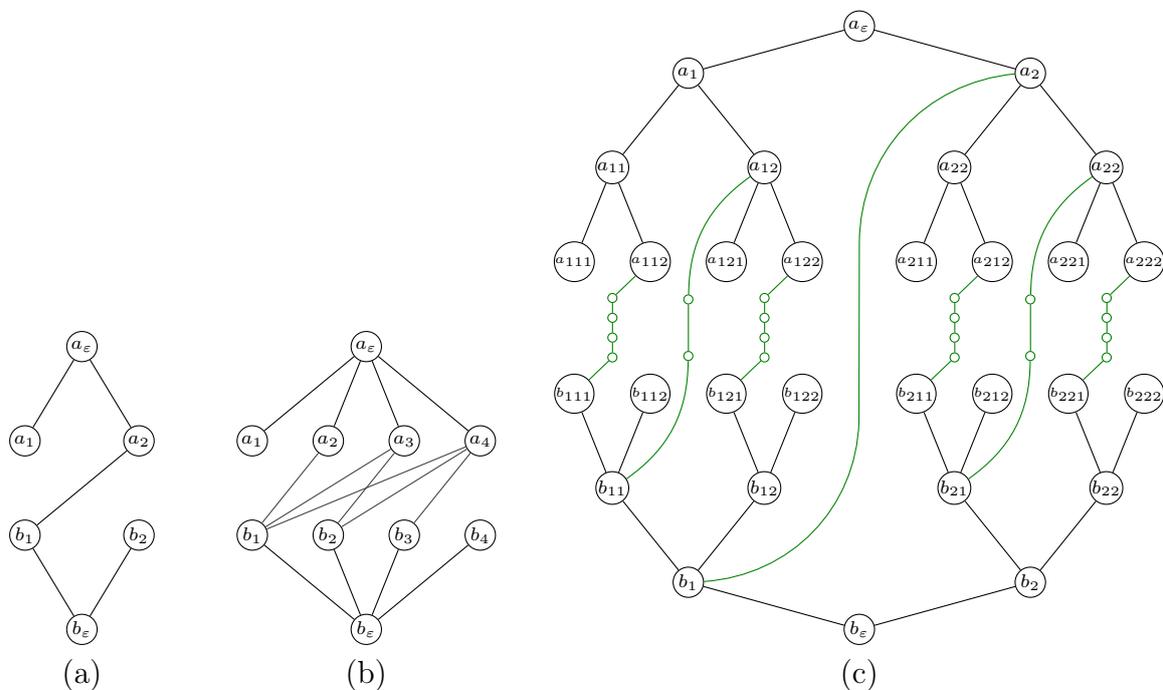


Figure 3: Graphs: (a) — $H_{2,1}$, (b) — $H_{4,1}$, (c) — $H_{2,3}$. ε denotes the empty word.

We now need to prove the required properties of $H_{k,h}$. Firstly, we prove that each vertex has at most $2k$ neighbors in the graph. Vertices belonging to neither \mathcal{A} nor \mathcal{B} lie on the paths connecting \mathcal{A} with \mathcal{B} , so their degrees are equal to 2. Let us now fix a vertex belonging to \mathcal{A} or \mathcal{B} . In its own tree, this vertex has at most k children and one parent. This vertex is also an endpoint of at most $k - 1$ paths connecting \mathcal{A} with \mathcal{B} ; the bound is satisfied with equality for vertices $a_{s \cdot k}$ and $b_{s \cdot 1}$ ($1, k \in \Sigma$, $s \in \Sigma^*$, $|s| \leq h - 1$). Hence, the total number of neighbors of each vertex of \mathcal{A} and \mathcal{B} is bounded by $2k$.

Secondly, each tree obviously has k^h leaves, since the leaves are parameterized by a word of length h over an alphabet of size k .

Finally, we need to find a half graph in this graph. Let (t_1, t_2, \dots, t_k) be the sequence of all words of length h over Σ , ordered lexicographically. The following lemma proves that the lexicographical ordering of the leaves of \mathcal{A} and \mathcal{B} produces a half graph:

Lemma 6. Fix two indices $i, j \in [1, k^h]$. The distance between b_{t_i} and a_{t_j} is equal to $2h - 1$ if $i < j$, and is greater than $2h - 1$ if $i \geq j$.

Proof. For simplicity, we can assume that each path of length ℓ ($\ell \geq 1$) connecting a vertex of \mathcal{A} with a vertex of \mathcal{B} is, in fact, a weighted edge of length ℓ connecting these two vertices.

We introduce the **level function** $\xi : V(\mathcal{A}) \cup V(\mathcal{B}) \rightarrow \mathbb{N}$, defined in the following way:

$$\xi(v) = \begin{cases} h - |s| & \text{if } v = a_s \text{ for } s \in \Sigma^*, \\ h + |s| - 1 & \text{if } v = b_s \text{ for } s \in \Sigma^*. \end{cases}$$

For a vertex v , we say that the value $\xi(v)$ is the **level** of vertex v .

We see that the levels of the vertices of \mathcal{A} range between 0 and h , while the levels of the vertices of \mathcal{B} range between $h - 1$ and $2h - 1$ (Figure 6).

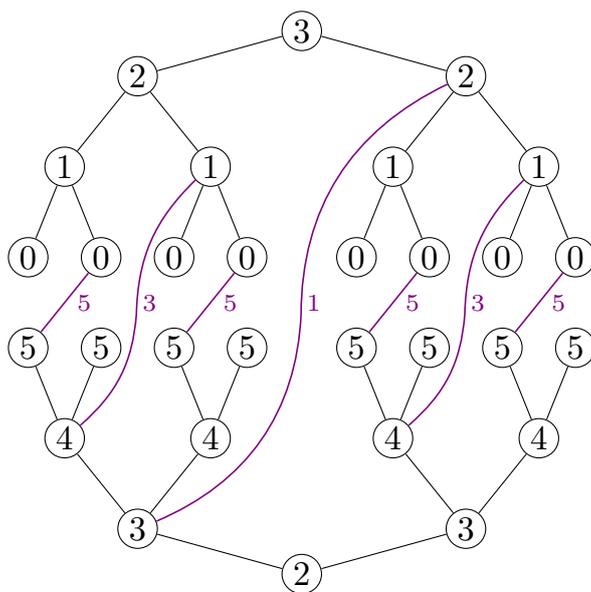


Figure 4: The levels of vertices of $H_{2,3}$. Each weighted edge resulting from a contracted path is colored violet and is labeled by its length.

We can now observe that:

- For every two vertices of \mathcal{A} connected by an edge — a_s and $a_{s.c}$ for some $s \in \Sigma^*$, $c \in \Sigma$ — we have that $\xi(a_s) = \xi(a_{s.c}) + 1$.
- For every two vertices of \mathcal{B} connected by an edge — b_s and $b_{s.c}$ — we have that $\xi(b_{s.c}) = \xi(b_s) + 1$.

- For every two vertices $a_{s,d} \in V(\mathcal{A})$, $b_{s,c} \in V(\mathcal{B})$ ($s \in \Sigma^*$, $c, d \in \Sigma$, $c < d$), connected by a weighted edge of length $2|s| + 1$, we have

$$\xi(b_{s,c}) - \xi(a_{s,d}) = (h + |s|) - (h - |s| - 1) = 2|s| + 1.$$

Hence, for each edge of weight x connecting two vertices u, v , we have that $x = |\xi(u) - \xi(v)|$. Therefore, for all pairs of vertices u, v , the distance between these vertices is bounded from below by $|\xi(u) - \xi(v)|$, and the bound is satisfied with equality only if there exists a path between u and v on which the values of ξ are monotone. We will call such paths **monotonous**.

Fix $i, j \in [1, k^h]$, and consider the vertices b_{t_i} and a_{t_j} . By the definition of ξ , we know that $\xi(a_{t_j}) = 0$ and $\xi(b_{t_i}) = 2h - 1$, hence the distance between these two vertices is at least $2h - 1$. We will now analyze when this distance is equal to $2h - 1$.

We see that the value of ξ only increases when: going up the tree \mathcal{A} (in the direction of the root of \mathcal{A}), taking a weighted edge from a vertex of \mathcal{A} to a vertex of \mathcal{B} (but not the other way around), and going down the tree \mathcal{B} (away from the root of \mathcal{B}). Hence, each path between b_{t_i} and a_{t_j} of length $2h - 1$ can only have such three segments, in this exact order, as it needs to be monotonous.

If $i < j$, then $t_i \prec t_j$, so the words t_i and t_j admit the following factorization:

$$t_i = p \cdot c \cdot s_i, \quad t_j = p \cdot d \cdot s_j$$

for $p, s_i, s_j \in \Sigma^*$, $c, d \in \Sigma$ and $c < d$.

Therefore, the path originating at a_{t_j} , going up \mathcal{A} to the vertex $a_{p,d}$, taking the weighted edge to $b_{p,c}$ connecting both trees (this is possible since $c < d$), and then going down \mathcal{B} to the destination b_{t_i} , is monotonous. Hence, $\text{dist}(b_{t_i}, a_{t_j}) = 2h - 1$.

Conversely, if a path from a_{t_j} to b_{t_i} has length $2h - 1$, then it is monotonous. Therefore, it must lead from a_{t_j} to an ancestor of a_{t_j} (we will name it a_x where x is a prefix of t_j), then take a weighted edge to an ancestor of b_{t_i} (we will call it a_y where y is a prefix of t_i), and then walk down the tree \mathcal{B} to b_{t_i} . By the construction of the tree, x and y must be both of the form $x = p \cdot d$, $y = p \cdot c$ where $p \in \Sigma^*$, $c, d \in \Sigma$, $d > c$. Since x and y are prefixes of t_j and t_i , respectively, this means that $t_j \succ t_i$. Hence, $j > i$.

Therefore, the path of length $2h - 1$ exists between a_{t_j} and b_{t_i} if and only if $i < j$; otherwise the shortest path between these vertices is longer. \square

Lemma 6 proves that the graph $H_{k,h}$ contains a half graph of order k^h . Since $k \geq 2$ and $h \geq 1$ were two arbitrary variables, this statement is satisfied for each choice of these variable.

Let us now finish the proof of Theorem 5. For each $\Delta \geq 4$ and odd $d \geq 1$ we can see that the graph $H_{\lfloor \frac{\Delta}{2} \rfloor, \frac{d+1}{2}}$ has its maximum degree bounded by $2 \lfloor \frac{\Delta}{2} \rfloor \leq \Delta$, and contains a distance- d half graph of order $\lfloor \frac{\Delta}{2} \rfloor^{\lceil \frac{d}{2} \rceil}$. Hence, the proof of the theorem is complete.

We note the following immediate conclusion from Theorem 5:

Corollary 7. *Let \mathcal{C}_Δ be the class of graphs with maximum degree bounded by $\Delta \geq 4$. Then the maximum order of a distance- d half graph ($d \geq 1$) in \mathcal{C}_Δ is bounded from below by $\Delta^{\Omega(d)}$.*

It turns out that the presented family of examples also witness as an exponential lower bound on the maximum order of a distance- d half graph in the class of planar graphs. This is proved explained by the following lemma:

Lemma 8. *For each $h \geq 1$, the graph $H_{2,h}$ is planar.*

Proof. For each $h \geq 1$, let $H'_{2,h}$ be the graph obtained from $H_{2,h}$ by contracting each path connecting two trees to a single edge. Obviously, $H'_{2,h}$ is planar if and only if $H_{2,h}$ is planar. Then, for $h \geq 2$, the graph $H'_{2,h}$ can be created from two copies of the graph $H'_{2,h-1}$ in the following way:

1. In the i -th copy of the graph ($i \in \{1, 2\}$), replace each vertex a_s ($s \in \{1, 2\}^*$) by $a_{i,s}$, and analogously each vertex b_s by $b_{i,s}$.
2. Create new roots of \mathcal{A} and \mathcal{B} — a_ε and b_ε , respectively.
3. Add the following edges: $a_\varepsilon a_1$, $a_\varepsilon a_2$, $b_\varepsilon b_1$, $b_\varepsilon b_2$, $b_1 a_2$.

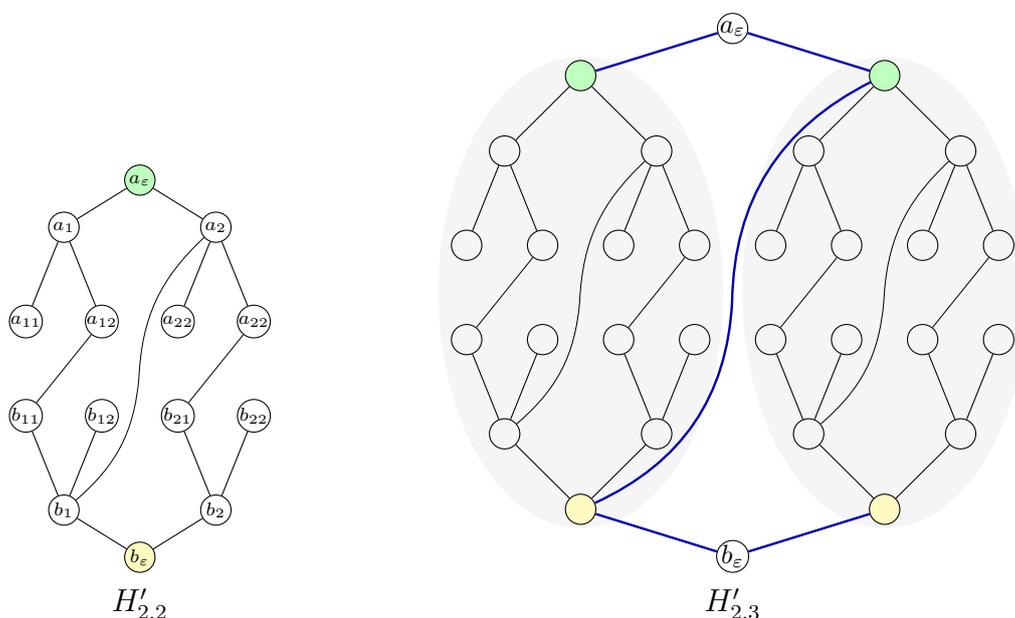


Figure 5: A step in the inductive process of creating $H'_{2,h}$ for each $h \geq 2$. Two copies of $H'_{2,2}$ (left; its roots are marked green and yellow) are put side by side, two new vertices are added, and five new edges (blue) are drawn, resulting in the graph $H'_{2,3}$ (right).

Note that $H'_{2,1} = H_{2,1}$ is planar (Figure 3(a)), and moreover, in the presented embedding, its vertices a_ε and b_ε lie on the outer face of the graph. These two properties can be maintained inductively: for $h \geq 2$, we take the embeddings of two copies of $H'_{2,h-1}$ satisfying these properties, put them side by side, and add the required edges, making sure that the graph remains planar and that the new roots of the trees are on the outer face of the graph (Figure 5). \square

We immediately infer the following:

Corollary 9. For every $d \geq 1$, the maximum order of a distance- d half graph ($d \geq 1$) in the class of planar graphs is bounded from below by $2^{\lceil \frac{d}{2} \rceil}$.

Proof. Assume d is odd. Then, by Lemma 8, the graph $H_{2, \frac{d+1}{2}}$ is planar. Also, it contains a distance- d half graph of order $2^{\frac{d+1}{2}} = 2^{\lceil \frac{d}{2} \rceil}$.

For even d , we construct a planar graph containing a distance- $(d-1)$ half graph of order $2^{\frac{d}{2}} = 2^{\lceil \frac{d}{2} \rceil}$, and then for each vertex a_i in this graph ($1 \leq i \leq 2^{\frac{d}{2}}$), we add an edge whose one endpoint is a_i , and another is a new vertex a'_i . We now can see that the vertices $a'_1, b_1, a'_2, b_2, \dots, a'_{2^{\frac{d}{2}}}, b_{2^{\frac{d}{2}}}$ form a distance- d half graph of required order. \square

3.2 Graphs with bounded pathwidth

In this section, we will find large half graphs in the class of graphs with bounded pathwidth. Specifically, we will prove the following theorem:

Theorem 10. For each $d \geq 1$ and $p \geq 0$, there exists a graph $P_{p,d}$ with pathwidth at most $p+2$ which contains a distance- $(4d-1)$ half graph of order $(2d+1)^p$.

Proof. Let $\ell_{p,d} := (2d+1)^p$. We define a slightly broader class of graphs $\{P_{p,d,k} \mid d \geq 1, p \geq 0, k \geq 0\}$ where $P_{p,d,k}$ has the following properties:

- The graph has pathwidth at most $p+2$.
- The graph contains a distance- $(4d-1)$ half graph $a_1, a_2, \dots, a_{\ell_{p,d}}, b_1, b_2, \dots, b_{\ell_{p,d}}$ in which every two vertices are at distance $2d$ or more from each other.
- Each of the vertices of the half graph has k vertex disjoint paths of length $3d$ attached to it (called *appendices*); in other words, after removing any of these vertices, there are at least k connected components which are paths with $3d$ vertices.

Thanks to the first two conditions, the graph $P_{p,d} := P_{p,d,0}$ will satisfy Theorem 10. The final condition is purely technical and will help us argue that pathwidths of the graphs constructed in the process described below do not grow too large.

The construction will be inductive over p . First, for $d \geq 1, k \geq 0$, we choose the graph $P_{0,d,k}$ to be the $3d$ -subdivision of two disjoint stars with k edges, with the distance- $(4d-1)$ half graph of order $\ell_{0,d} = 1$ given by the centers of both stars. It can be easily verified that this graph satisfies the requirements; in particular, $\text{pw}(P_{0,d,k}) \leq 2$.

We shall now describe how to construct $P_{p,d,k}$ for $p, d \geq 1$ and $k \geq 0$:

1. Take $2d+1$ disjoint copies of $P_{p-1,d,k+1}$, and renumber the vertices of the half graph in each of the copies, ordering all $2d+1$ half graphs one after another. Formally, the labels of the vertices a_i and b_i of the half graph in the j -th disjoint copy of the graph are changed to $a_{(j-1)\ell_{p-1,d}+i}$ and $b_{(j-1)\ell_{p-1,d}+i}$, respectively. Since $\ell_{p,d} = (2d+1)\ell_{p-1,d}$, the graph contains vertices a_i, b_i for each $i \in [1, \ell_{p,d}]$. The following steps of the construction will ensure that these vertices form a distance- $(4d-1)$ half graph, provided that $P_{p-1,d,k+1}$ satisfies the required properties.
2. Add a new vertex r to the graph.

- For each vertex b_i in the j -th disjoint copy of the graph, $j \in [1, 2d + 1]$, we connect b_i with r with a path of length $d + (j - 1)$. In order to do that, we use up one of $k + 1$ appendices attached to b_i by connecting r with the $(d + j - 2)$ -nd vertex of the appendix with an edge (Figure 6); the part of the appendix after this vertex can then be safely discarded.
- We apply a similar procedure to the vertices a_i of the half graph, only that the vertex of the half graph in the j -th disjoint copy of the graph is connected with r with a path of length $3d - (j - 1)$. This finishes the construction of $P_{p,d,k}$.

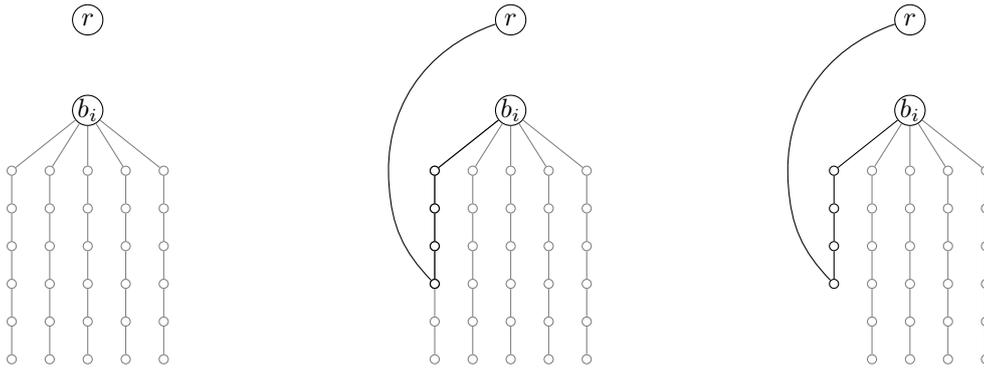


Figure 6: The process of turning an appendix attached to b_i into a path of length 5 connecting b_i with r .

Obviously, each vertex a_i and b_j in $P_{p,d,k}$ has k remaining appendices. We also note that $P_{p,d,k} - \{r\}$ is a subgraph of a graph formed from $2d + 1$ disjoint copies of $P_{p-1,d,k+1}$. Hence,

$$\text{pw}(P_{p,d,k}) \leq \text{pw}(P_{p,d,k} - \{r\}) + 1 \leq \text{pw}(P_{p-1,d,k+1}) + 1 \leq p + 2.$$

We also argue that $a_1, a_2, \dots, a_{\ell_{p,d}}, b_1, b_2, \dots, b_{\ell_{p,d}}$ is a distance- $(4d - 1)$ half graph satisfying the distance requirements of $P_{p,d,k}$. In the construction, multiple disjoint copies of $P_{p-1,d,k+1}$ were connected by vertex disjoint paths to a new vertex r , where each path has length at least d and at most $3d$. Hence, the shortest path between any two vertices of the new half graph is either fully contained in a copy of $P_{p-1,d,k+1}$ (and it must not be shorter than $2d$ by the inductive assumption), or it passes through r and thus it follows two paths of length at least d , so its length is at least $2d$ as well. From this property, we infer that if the shortest path between any two vertices of the half graph has length at most $4d - 1$, then this path must avoid all other vertices of the half graph.

Let us now consider whether there exists a path between two vertices a_i and b_j of the half graph of length $4d - 1$. Again, this path must either be fully contained in a single copy of $P_{p-1,d,k+1}$ (and this condition is only satisfied if $i > j$), or it passes through r . In the latter case, this path must consist of a direct connection from a_i to r , and then a direct connection from r to b_j . Assuming that a_i is in the x -th disjoint copy of $P_{p-1,d,k+1}$, and b_j is in the y -th copy of the graph ($x, y \in [1, 2d + 1]$), we infer that the length of that path is exactly

$$[d + (y - 1)] + [3d - (x - 1)] = 4d + (y - x).$$

Hence, such a path exists if and only if $x > y$; that is, only if $i > j$ and both a_i and b_j come from the different copies of $P_{p-1,d,k+1}$.

We conclude that $\text{dist}(a_i, b_j) \leq 4d - 1$ if and only if $i > j$. Therefore, all the required conditions are satisfied, so $P_{p,d,k}$ is defined correctly. This finishes the proof of the lower bound. \square

On an end note, we remark the following asymptotic version of Theorem 10:

Corollary 11. *For each pair of integers $p \geq 2$, $d \geq 4$, the maximum order of a distance- d half graph in the class of graphs with pathwidth bounded by p is bounded from below by $\Theta(d)^{p-O(1)}$.*

3.3 Graphs with bounded treewidth and minor-free graphs

In this section, we will construct a family of graphs with small treewidth containing large half graphs. Since the graphs with bounded treewidth also avoid large cliques as minors, this family will also serve as examples of graphs avoiding relatively small cliques as minors, but still containing large half graphs.

To start with, we shall introduce a variant of a tree decomposition of a graph G containing a distance- $2d$ half graph.

Definition 12. A rooted tree decomposition \mathcal{T} of a graph G is a **distance- d pairing decomposition** if all the following conditions hold:

- The root of \mathcal{T} is a bag containing exactly two elements, called A and B .
- Each leaf of \mathcal{T} is a bag containing exactly two elements, called a_i and b_i , where $i \in \{1, 2, \dots, \ell\}$ and ℓ is the number of leaves in \mathcal{T} .
- All vertices $A, B, a_1, b_1, a_2, b_2, \dots, a_\ell, b_\ell$ are pairwise different.
- For each pair of integers i, j such that $1 \leq i < j \leq \ell$, we have that $\text{dist}(b_i, a_j) = 2d$.
- For each pair of integers i, j such that $1 \leq j \leq i \leq \ell$, we have that $\text{dist}(b_i, a_j) = 2d + 1$.

For such a tree decomposition, we define two objects: its root $\text{Root}(\mathcal{T}) = (A, B)$ and the sequence of vertices in its leaves $\text{Leaves}(\mathcal{T}) = (a_1, b_1, a_2, b_2, \dots, a_\ell, b_\ell)$. We will refer to the value ℓ as the **order** of \mathcal{T} . We also define the **width** of the tree decomposition in the standard way.

We remark that Definition 12 implies that the set of vertices $a_1, a_2, \dots, a_\ell, b_1, b_2, \dots, b_\ell$ forms a distance- $2d$ half graph of order ℓ in G . However, the conditions required here are even stronger as the definition of a half graph only requires that $\text{dist}(b_i, a_j) \leq 2d$ for each pair of integers $i < j$, and $\text{dist}(b_i, a_j) > 2d$ for $i > j$. This strengthening will be essential in the proof of the correctness of the construction.

We also note that \mathcal{T} — in some way — “exposes” the vertices of the underlying half graph in the form of the leaf bags of the tree decomposition ($\text{Leaves}(\mathcal{T})$), and two other vertices in the form of the root bag of the decomposition ($\text{Root}(\mathcal{T})$). These outer bags of the tree decomposition are “sockets” which can connect to the “sockets” of pairing decompositions of other graphs, allowing a recursive construction of large graphs.

Example 13. We consider the graph G presented in Figure 7(a) with the tree decomposition \mathcal{T} depicted in Figure 7(b). This decomposition is distance-1 pairing and has order 2 and width 3. The table in Figure 7(c) shows the distances between a_i and b_j for each pair of indices i, j . We can see that $\text{dist}(b_1, a_2) = 2$ and all the remaining distances are equal to 3. Moreover, $\text{Root}(\mathcal{T}) = (A, B)$ and $\text{Leaves}(\mathcal{T}) = (a_1, b_1, a_2, b_2)$.

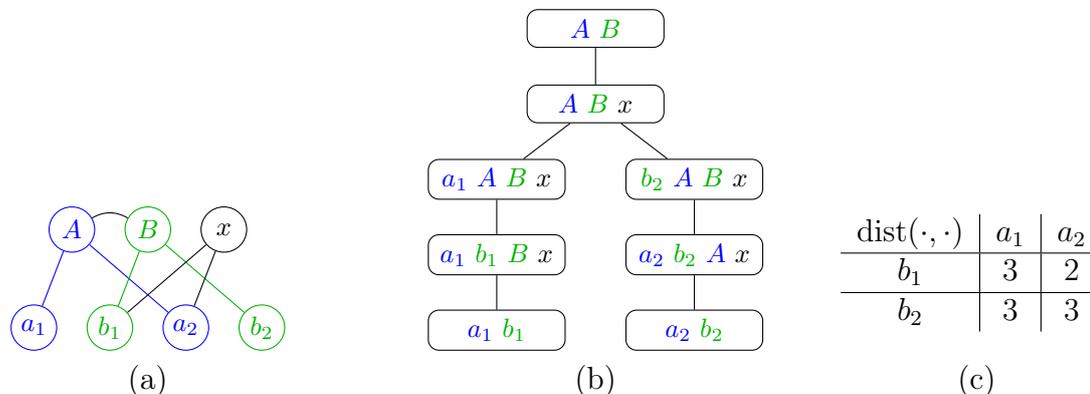


Figure 7: (a) — the graph G in Example 13; (b) — its tree decomposition; (c) — its distance matrix.

In the construction of our family of graphs, we will need to use a refinement of Definition 12. Intuitively, we want to expose the root vertices of the decomposition in an efficient way. We will do it by creating a stronger variant of the pairing decomposition in which one of the root vertices is adjacent to all the vertices on one side of the half graph, while the other root vertex is adjacent to the other side.

Definition 14. Consider an arbitrary undirected graph G with a distance- d pairing decomposition \mathcal{T} where $\text{Root}(\mathcal{T}) = (A, B)$ and $\text{Leaves}(\mathcal{T}) = (a_1, b_1, \dots, a_\ell, b_\ell)$. We will say that \mathcal{T} is **neighboring** if for each $i \in \{1, 2, \dots, \ell\}$, we have that:

$$\text{dist}(a_i, A) = \text{dist}(b_i, B) = 1, \quad \text{dist}(a_i, B) \geq 2d + 1, \quad \text{dist}(b_i, A) \geq 2d + 1.$$

Moreover, A must be adjacent only to vertices a_1, \dots, a_ℓ , and B must be adjacent only to vertices b_1, \dots, b_ℓ .

We can verify that the decomposition from Example 13 is not neighboring since the vertices A and B are connected by an edge.

We will now present a few lemmas. Lemma 15 will allow us to turn an arbitrary decomposition into a neighboring decomposition by incurring a small cost on the width of the decomposition. Lemma 16 will be a method of combining two pairing decompositions — a neighboring decomposition of width t and order ℓ_1 , and a pairing decomposition of width $t - 2$ and order ℓ_2 , into a pairing decomposition of order $\ell_1\ell_2$. Then, we will combine both lemmas in order to create graphs with enormous pairing decompositions.

Lemma 15. For every undirected graph G with a distance- d pairing decomposition \mathcal{T} of width t and order ℓ , there exists a graph G' with a neighboring distance- d pairing decomposition \mathcal{T}' of width $t + 2$ and order ℓ .

Proof. Let $\text{Leaves}(\mathcal{T}) = (a_1, b_1, \dots, a_\ell, b_\ell)$.

We construct G' from G by creating two new vertices A_{new} and B_{new} , connecting A_{new} directly by an edge to each of the vertices a_i for $i \in [1, \ell]$, and analogously connecting B_{new} to each of the vertices b_i .

We will now construct a rooted tree decomposition \mathcal{T}' from \mathcal{T} in the following way

1. For each leaf bag $L_i \in V(\mathcal{T})$ containing vertices a_i and b_i , we clone it (without cloning the vertices in the bag), and connect it with L_i in \mathcal{T} . This way, the clone becomes a leaf in the new decomposition instead of L_i .
2. We create a new bag $(A_{\text{new}}, B_{\text{new}})$, which we directly connect to the root of \mathcal{T} .
3. To each bag of \mathcal{T} which is neither the root nor a leaf of the new tree decomposition, we add both A_{new} and B_{new} .

We can easily see that \mathcal{T}' is a tree decomposition of G' of width $t + 2$.

We will now show that (G', \mathcal{L}') is a neighboring decomposition by deducing all the required equalities and inequalities between the distances. Obviously, $\text{dist}(a_i, A_{\text{new}}) = \text{dist}(b_i, B_{\text{new}}) = 1$ for each $i \in [1, \ell]$. Furthermore, by considering all neighbors of B_{new} , we can see that for each $i \in [1, \ell]$, we have that

$$\text{dist}(a_i, B_{\text{new}}) = \min_{j \in \{1, \dots, \ell\}} (\text{dist}(a_i, b_j) + 1) \geq 2d + 1.$$

Analogously, $\text{dist}(b_i, A_{\text{new}}) \geq 2d + 1$ for every $i \in [1, \ell]$.

We can now prove that no path connecting b_i and a_j which has length $2d$ or less, for any choice of indices $i, j \in [1, \ell]$, can pass through A_{new} :

$$\text{dist}(b_i, A_{\text{new}}) \geq 2d + 1 \quad \Rightarrow \quad \text{dist}(b_i, A_{\text{new}}) + \text{dist}(A_{\text{new}}, a_j) > 2d + 1.$$

Analogously, no such path can pass through B_{new} .

If $i < j$, then $\text{dist}_G(b_i, a_j) = 2d$ holds in G . Since A_{new} and B_{new} were the only vertices added to G in the construction of G' , and no path with length $2d$ or less can pass through these vertices, this means that we also have $\text{dist}_{G'}(b_i, a_j) = 2d$. Analogously, if $i \geq j$, then $\text{dist}_G(b_i, a_j) = 2d + 1$ in also implies $\text{dist}_{G'}(b_i, a_j) = 2d + 1$. \square

Following the notation from the statement of Lemma 15, we define

$$\text{MakeNeighboring}(G, \mathcal{T}) := (G', \mathcal{T}').$$

Lemma 16. *For every $d \geq 2$, $t \geq 1$, $\ell_1, \ell_2 \geq 2$, and for every pair of graphs G, H where:*

- G has a distance- $(d - 1)$ pairing decomposition \mathcal{T}_G of order ℓ_1 and width t ,
- H has a neighboring distance- d pairing decomposition \mathcal{T}_H of order ℓ_2 and width t ,

there exists a graph with a distance- d pairing decomposition of order $\ell_1 \ell_2$ and width t .

Proof. We fix the graph G with its pairing decomposition \mathcal{T}_G , and let $\text{Root}(\mathcal{T}_G) = (A, B)$, $\text{Leaves}(\mathcal{T}_G) = (A_1, B_1, A_2, B_2, \dots, A_{\ell_1}, B_{\ell_1})$. We also create ℓ_1 disjoint copies of the graph H together with its tree decomposition \mathcal{T}_H . For each $i \in [1, \ell_1]$, we define the graph H_i as the i -th copy of H . In its tree decomposition \mathcal{T}_{H_i} , we name the vertices of the root bag and the leaf bags:

$$\begin{aligned} \text{Root}(\mathcal{T}_{H_i}) &= (\alpha_i, \beta_i), \\ \text{Leaves}(\mathcal{T}_{H_i}) &= (a_{(i-1)\ell_2+1}, b_{(i-1)\ell_2+1}, a_{(i-1)\ell_2+2}, b_{(i-1)\ell_2+2}, \dots, a_{i\ell_2}, b_{i\ell_2}). \end{aligned}$$

We remark that the vertices $A_1, B_1, A_2, B_2, \dots, A_{\ell_1}, B_{\ell_1}$ form a half graph in G . Moreover, for each $i \in [1, \ell_1]$, the set of vertices a_j, b_j for all $j \in [(i-1)\ell_2 + 1, i\ell_2]$ forms a half graph in H_i . We also note that in the sequence $(a_1, b_1, a_2, b_2, \dots, a_{\ell_1\ell_2}, b_{\ell_1\ell_2})$, the vertices of H_1 form a prefix, followed by the vertices of H_2 , then H_3 , and so on, and a suffix of this sequence is the vertices of H_{ℓ_1} .

We create a new undirected graph U , together with its tree decomposition \mathcal{T}_U , in the following way (Figure 8):

1. We take the disjoint union of the graphs $G, H_1, \dots, H_{\ell_1}$, together with the forest created by the disjoint union of the corresponding tree decompositions $\mathcal{T}_G, \mathcal{T}_{H_1}, \mathcal{T}_{H_2}, \dots, \mathcal{T}_{H_{\ell_1}}$.
2. For each $i \in \{1, 2, \dots, \ell_1\}$, we fuse the vertices A_i and α_i , and the vertices B_i and β_i .
3. In the forest of tree decompositions, for each $i \in \{1, 2, \dots, \ell_1\}$, we fuse the i -th leaf bag of \mathcal{T}_G (containing vertices A_i and B_i) with the root bag of \mathcal{T}_{H_i} (containing α_i and β_i). This step is possible since $A_i = \alpha_i$ and $B_i = \beta_i$. We call the resulting tree \mathcal{T}_U .

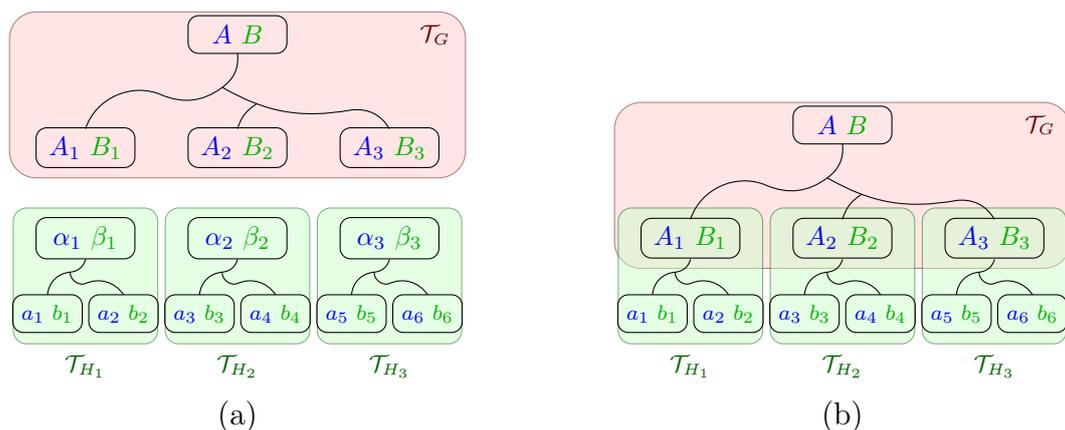


Figure 8: (a) — tree decompositions after Step 1 of the algorithm; all the bags apart from the root and leaves have been omitted.

(b) — the resulting tree decomposition \mathcal{T}_U .

It is straightforward to see that \mathcal{T}_U is a tree decomposition of U , and that its width remains equal to t .

In U , we can identify the subgraphs $G, H_1, \dots, H_{\ell_1}$, originating from the disjoint union in Step 1. We stress that for each $i \in [1, \ell_1]$, subgraphs G and H_i share two vertices — A_i and B_i . These two vertices form a bag separating \mathcal{T}_G from \mathcal{T}_{H_i} . This has a couple of consequences, directly implied by the properties of tree decompositions:

- If a simple path P connects a vertex from H_i with a vertex outside of H_i , then P must pass through one of the vertices A_i, B_i (possibly both);
- If a simple path P has both endpoints in H_i , but contains a vertex outside of H_i , then P must pass through both A_i and B_i .

We now proceed to proving all the distance equalities and inequalities required by the definition of a distance- d pairing decomposition.

Claim 17. *For each $k \in [1, \ell_1]$, the vertices A_k and B_k are at distance $2d + 2$ in H_k .*

Proof. The set of neighbors of A_k in H_k is exactly $\{a_j \mid j \in [(k-1)\ell_2 + 1, k\ell_2]\}$, while the set of neighbors of B_k in H_k is $\{b_i \mid i \in [(k-1)\ell_2 + 1, k\ell_2]\}$. Hence, the second vertex on the shortest path from A_k to B_k in H_k is a_j for some $j \in [(k-1)\ell_2 + 1, k\ell_2]$, while the penultimate vertex on this path is b_i for some $i \in [(k-1)\ell_2 + 1, k\ell_2]$.

Since the vertices a_i, b_i for $i \in [(k-1)\ell_2 + 1, k\ell_2]$ are in the distance- d pairing decomposition \mathcal{T}_{H_k} , we get that $\text{dist}_{H_k}(b_i, a_j) = 2d$ if $i < j$; otherwise, this distance is larger than $2d$. As $\ell_2 \geq 2$, the vertices $b_{k\ell_2-1}$ and $a_{k\ell_2}$ both belong to H_k , and the distance between them is equal to $2d$. Hence,

$$\text{dist}_{H_k}(A_k, B_k) = 2 + \min_{i, j \in [(k-1)\ell_2 + 1, k\ell_2]} \text{dist}_{H_k}(b_i, a_j) = 2 + \text{dist}_{H_k}(b_{k\ell_2-1}, a_{k\ell_2}) = 2d + 2.$$

┘

Claim 18. *For each pair of vertices $u, v \in V(G)$, we have that $\text{dist}_U(u, v) = \text{dist}_G(u, v)$.*

Proof. We first note that for each $k \in [1, \ell_1]$, the vertices A_k and B_k are at distance $2d - 1$ in G ; this is because \mathcal{T}_G is a distance- $(d - 1)$ pairing decomposition of G .

Let us fix two vertices $u, v \in V(G)$ and assume that the shortest path in U between them passes through a vertex $x \notin V(G)$. Then, $x \in V(H_k)$ for some $k \in [1, \ell_1]$. Hence, the path enters the subgraph H_k through one of the vertices A_k, B_k , passes through x , and then leaves this subgraph through the other of these vertices. This means that this shortest path contains a subpath connecting A_k and B_k , which is fully contained within H_k . By Claim 17, this subpath has length at least $2d + 2$, and therefore it can be replaced by the shortest path in G connecting A_k and B_k , which has length $2d - 1$. Therefore, the considered path was not a shortest path in the first place — a contradiction.

Hence, the shortest path between u and v lies fully within G and thus $\text{dist}_U(u, v) = \text{dist}_G(u, v)$. ┘

Claim 18 significantly simplifies the structure of the distances between the vertices of U ; no shortest path between a pair of vertices in U can enter a subgraph H_k for some $k \in [1, \ell_1]$ and then leave it, since there exists a strictly shorter path between these two vertices omitting this subgraph.

Also, as a corollary of Claim 18, we note the following: for each $k \in [1, \ell_1]$, we have that $\text{dist}_U(A_k, B_k) = \text{dist}_G(A_k, B_k) = 2d - 1$.

We will now prove a claim that will directly imply that \mathcal{T}_U is a distance- d pairing decomposition of U .

Claim 19. *For each pair of indices $i, j \in [1, \ell_1 \ell_2]$, we have that*

$$\text{dist}_U(b_i, a_j) = \begin{cases} 2d & \text{if } i < j, \\ 2d + 1 & \text{if } i \geq j. \end{cases}$$

Proof. Firstly, let us assume that b_i and a_j are located in the same subgraph H_k of U for some $k \in [1, \ell_1]$; that is, $i, j \in [(k - 1)\ell_2 + 1, k\ell_2]$. The shortest path in U between these two vertices will have one of the following shapes:

- The path is fully contained within H_k . In this case,

$$\text{dist}_U(b_i, a_j) = \text{dist}_{H_k}(b_i, a_j) = \begin{cases} 2d & \text{if } i < j, \\ 2d + 1 & \text{if } i \geq j, \end{cases}$$

as required.

- The path leaves H_k and then reenters this subgraph. This would, however, require the path to contain a path between A_k and B_k as a subpath. Since $\text{dist}_U(A_k, B_k) = 2d - 1$, and neither A_k nor B_k can be an endpoint of the considered path, we get that

$$\text{dist}_U(b_i, a_j) \geq 2 + \text{dist}_U(A_k, B_k) \geq 2d + 1.$$

However, there exists a path of equal or shorter length contained within H_k .

Therefore, there is a shortest path between b_i and a_j that is fully contained within H_k , which completes the proof in this case.

Let us now assume that $a_i, b_i \in V(H_s)$ and $a_j, b_j \in V(H_t)$ for some $s, t \in [1, \ell_1]$, $s < t$, and let us compute the distances $\text{dist}_U(b_i, a_j)$ and $\text{dist}_U(b_j, a_i)$. By the construction of U , we infer that $i < j$.

By our considerations above, we deduce that the shortest path from b_i to a_j originates in H_s , then enters G through either of the vertices A_s, B_s , then enters H_t through either of the vertices A_t, B_t , and eventually terminates in H_t (Figure 9). The same applies to the shortest path from b_j to a_i .

We will split this shortest path into three parts: the first part, terminating at either of the vertices A_s, B_s ; the second part, taking off from where the first path finished, and continuing to either of the vertices A_t, B_t ; and the final part, continuing until the reaching a_j or b_j . If there are multiple possible partitions, we pick the split minimizing the length of the first and the third part. This ensures us that the first path is fully contained within H_s , and the final part is fully contained within H_t . Since $A_s, B_s, A_t, B_t \in V(G)$, Claim 18 asserts that the second part of the path is also fully contained within G .

Since \mathcal{T}_{H_s} is a neighboring tree decomposition of H_s , we figure that

$$\text{dist}_{H_s}(a_i, A_s) = \text{dist}_{H_s}(b_i, B_s) = 1, \quad \text{dist}_{H_s}(a_i, B_s) \geq 2d + 1, \quad \text{dist}_{H_s}(b_i, A_s) \geq 2d + 1.$$

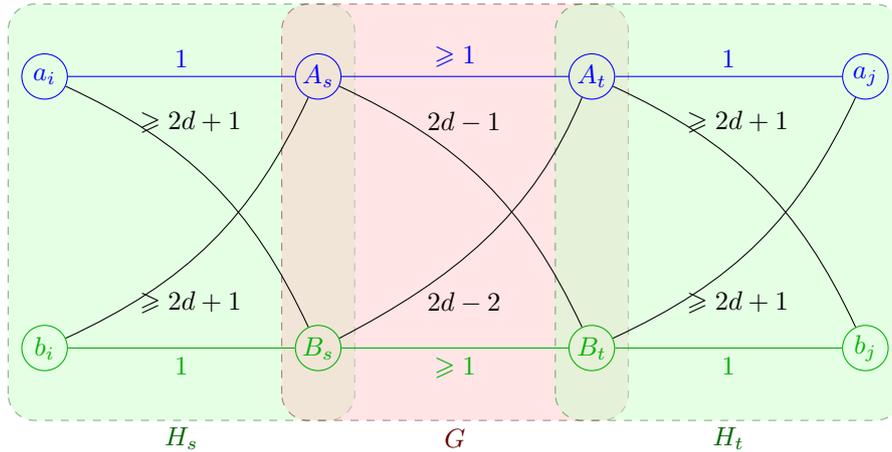


Figure 9: The important vertices on the shortest path from b_i to a_j , and on the shortest path from b_j to a_j . The edges are labeled by the length of the shortest path between the connected vertices or the lower bound on this length.

Analogously, in H_t :

$$\text{dist}_{H_t}(a_j, A_t) = \text{dist}_{H_t}(b_j, B_t) = 1, \quad \text{dist}_{H_t}(a_j, B_t) \geq 2d + 1, \quad \text{dist}_{H_t}(b_j, A_t) \geq 2d + 1.$$

Finally, in G , since $A_s \neq A_t$, we get that $\text{dist}_G(A_s, A_t) \geq 1$; similarly, $\text{dist}_G(B_s, B_t) \geq 1$. By the definition of G (which has a distance- $(d - 1)$ pairing decomposition) and the fact that $s < t$, we infer that $\text{dist}_G(B_s, A_t) = 2d - 2$ and $\text{dist}_G(B_t, A_s) = 2d - 1$.

By verifying all possibilities, we can now conclude that (Figure 9):

$$\begin{aligned} \text{dist}_U(b_i, a_j) &= \min_{X_s \in \{A_s, B_s\}} \min_{X_t \in \{A_t, B_t\}} [\text{dist}_{H_s}(b_i, X_s) + \text{dist}_G(X_s, X_t) + \text{dist}_{H_t}(X_t, a_j)] = \\ &= \text{dist}_{H_s}(b_i, B_s) + \text{dist}_G(B_s, A_t) + \text{dist}_{H_t}(A_t, a_j) = 2d \end{aligned}$$

and

$$\begin{aligned} \text{dist}_U(a_i, b_j) &= \min_{X_s \in \{A_s, B_s\}} \min_{X_t \in \{A_t, B_t\}} [\text{dist}_{H_s}(a_i, X_s) + \text{dist}_G(X_s, X_t) + \text{dist}_{H_t}(X_t, b_j)] = \\ &= \text{dist}_{H_s}(a_i, A_s) + \text{dist}_G(A_s, B_t) + \text{dist}_{H_t}(B_t, b_j) = 2d + 1. \end{aligned}$$

Hence, in this case, $i < j$ implies $\text{dist}(b_i, a_j) = 2d$ and $\text{dist}(b_j, a_i) = 2d + 1$. Therefore, the statement of the claim holds in this case as well.

Since the statement of the claim is true for each choice of a pair of indices i, j — no matter whether b_i and a_j are in the same subgraph H_k for some k or not — the proof is complete. \lrcorner

Claim 19 directly proves the lemma. \square

Following the notation from Lemma 16, we define

$$\text{CombineNeighboring}((G, \mathcal{T}_G), (H, \mathcal{T}_H)) := (U, \mathcal{T}_U).$$

Thus, `CombineNeighboring` is a function that takes:

- a graph G , together with its distance- $(d - 1)$ pairing decomposition \mathcal{T}_G of order ℓ_1 and width t ; and
- a graph H , together with its neighboring distance- d pairing decomposition \mathcal{T}_H of order ℓ_2 and width t ;

and returns a graph U and its distance- d decomposition \mathcal{T}_U of order $\ell_1\ell_2$ and width t .

We can now combine Lemmas 15 and 16:

Corollary 20. *For every $d \geq 2$, $t \geq 3$, $\ell_1, \ell_2 \geq 2$, and every pair of graphs G, H where:*

- G has a distance- $(d - 1)$ pairing decomposition \mathcal{T}_G of order ℓ_1 and width t , and
- H has a distance- d pairing decomposition \mathcal{T}_H of order ℓ_2 and width $t - 2$,

there exists a graph with a distance- d pairing decomposition of order $\ell_1\ell_2$ and width t .

Proof. We claim that the graph U and its pairing decomposition \mathcal{T}_U , defined as follows:

$$(U, \mathcal{T}_U) := \text{CombineNeighboring}((G, \mathcal{T}_G), \text{MakeNeighboring}(H, \mathcal{T}_H))$$

satisfies the statement of the lemma.

By Lemma 15, $\text{MakeNeighboring}(H, \mathcal{T}_H)$ is a graph with a distance- d neighboring pairing decomposition of order ℓ_2 and width t . By applying Lemma 16 on G and the newly produced graph, we create a graph U with a distance- d pairing decomposition \mathcal{T}_U of order $\ell_1\ell_2$ and width t . \square

Using the notation from the statement of Corollary 20, we set

$$\text{Combine}((G, \mathcal{T}_G), (H, \mathcal{T}_H)) := (U, \mathcal{T}_U).$$

Using Corollary 20, we will now produce huge pairing decompositions. Precisely, for each pair of integers $d \geq 1$, $k \geq 1$, we will define an undirected graph $G_{d,k}$, together with its distance- d pairing decomposition $\mathcal{T}_{d,k}$ of width $2k + 1$.

Since the tree decomposition of the graph in Example 13 is distance-1 pairing and has width 3, we shall take it as $\mathcal{T}_{1,1}$, and the corresponding graph as $G_{1,1}$. The decomposition $\mathcal{T}_{1,1}$ has order 2.

For $d \geq 2$, we create the graph $G_{d,1}$ and its distance- d decomposition $\mathcal{T}_{d,1}$ from $G_{1,1}$ and $\mathcal{T}_{1,1}$ in the following way: for each vertex $v \in \{a_1, b_1, a_2, b_2\}$, we append to the graph a path of length $d - 1$ whose one endpoint is v , and the other endpoint is a new vertex $v' \in \{a'_1, b'_1, a'_2, b'_2\}$. We can easily create a distance- d pairing decomposition $\mathcal{T}_{d,1}$ of $G_{d,1}$ of width 3 (Figure 10). We remark that $\mathcal{T}_{d,1}$ also has order 2.

Also, for $k \geq 2$, we create the graph $G_{1,k}$ and its distance-1 pairing decomposition $\mathcal{T}_{1,k}$ from $G_{1,1}$ and $\mathcal{T}_{1,1}$ by adding $2(k - 1)$ isolated vertices to $G_{1,1}$, and appending each new vertex to each bag of the decomposition $\mathcal{T}_{1,1}$. The resulting decomposition is obviously a distance-1 pairing decomposition of width $2k + 1$ and order 2.

Eventually, for $d \geq 2$, $k \geq 2$, we create the graph $G_{d,k}$ and the corresponding decomposition $\mathcal{T}_{d,k}$ by applying **Combine**:

$$(G_{d,k}, \mathcal{T}_{d,k}) := \text{Combine}((G_{d-1,k}, \mathcal{T}_{d-1,k}), (G_{d,k-1}, \mathcal{T}_{d,k-1})).$$

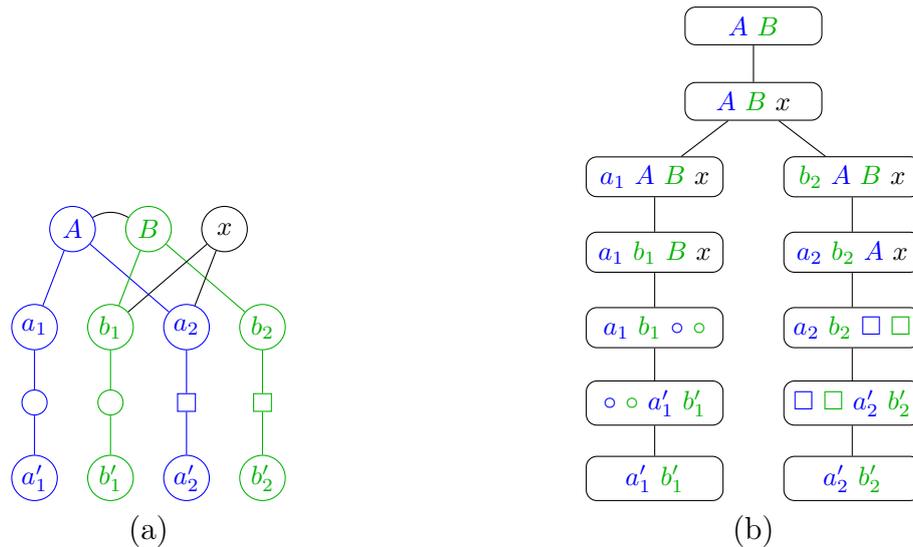


Figure 10: (a) — the graph $G_{3,1}$; (b) — an example tree decomposition $T_{3,1}$.

We can easily verify that the input to this function complies with Corollary 20, and leads to the construction of a graph with a distance- d pairing decomposition of width $2k + 1$, whose order is equal to the product of the orders of $\mathcal{T}_{d-1,k}$ and $\mathcal{T}_{d,k-1}$.

For $d, k \geq 1$ we set $\ell_{d,k}$ as the order of $\mathcal{T}_{d,k}$.

Lemma 21. *For each $d, k \geq 1$, the following holds:*

$$\ell_{d,k} = 2^{\binom{d+k-2}{k-1}}.$$

Proof. The formula can be directly verified for $d = 1$ and for $k = 1$. For $d, k \geq 2$, we know that the order of $\mathcal{T}_{d,k}$ is the product of the orders of $\mathcal{T}_{d-1,k}$ and $\mathcal{T}_{d,k-1}$. Therefore, by applying a simple induction on d and k , we compute that

$$\log_2 \ell_{d,k} = \log_2 \ell_{d-1,k} + \log_2 \ell_{d,k-1} = \binom{d+k-3}{k-1} + \binom{d+k-3}{k-2} = \binom{d+k-2}{k-1}.$$

□

Lemma 21 immediately leads to the following conclusion:

Theorem 22. *For every even $d \geq 2$ and odd $t \geq 3$, there exists a graph with treewidth not exceeding t that contains a distance- d half graph of order*

$$2^{\binom{\frac{1}{2}(d+t-5)}{\frac{1}{2}(t-3)}}.$$

Proof. We consider the graph $G_{\frac{d}{2}, \frac{t-1}{2}}$ together with its distance- $\frac{d}{2}$ pairing decomposition $\mathcal{T}_{\frac{d}{2}, \frac{t-1}{2}}$. The decomposition has order $\ell_{\frac{d}{2}, \frac{t-1}{2}}$ (Lemma 21), which matches the formula in the statement of the theorem; and width t , which ensures that the graph has treewidth not exceeding t . Eventually, by the properties of pairing decompositions, $\text{Leaves}(\mathcal{T}_{\frac{d}{2}, \frac{t-1}{2}})$ is a distance- d half graph of order $\ell_{\frac{d}{2}, \frac{t-1}{2}}$. □

We also remark the asymptotic version of this theorem.

Corollary 23. *Fix $t \geq 3$. Then for $d \geq 2$, the maximum order of a distance- d half graph in the class of graphs with treewidth at most t is bounded from below by $2^{d^{\Omega(t)}}$.*

This discovery immediately leads to a lower bound in the class of K_t -minor-free graphs:

Theorem 24. *For every even $d \geq 2$ and odd $t \geq 5$, there exists a K_t -minor-free graph containing a distance- d half graph of order*

$$2^{\binom{\frac{1}{2}(d+t-7)}{\frac{1}{2}(t-5)}}.$$

Proof. Theorem 22 asserts the existence of a graph with treewidth not exceeding $t - 2$ containing a distance- d half graph of order $2^{\binom{\frac{1}{2}(d+t-7)}{\frac{1}{2}(t-5)}}$. By Theorem 4, this graph is K_t -minor-free. \square

We immediately infer the following asymptotic result:

Corollary 25. *Fix $t \geq 5$. Then for $d \geq 2$, the maximum order of a distance- d half graph in the class of K_t -minor-free graphs is bounded from below by $2^{d^{\Omega(t)}}$.*

4 Polynomial bound on neighborhood complexity in planar graphs

In this section, we prove that the neighborhood complexity of any vertex set S in planar graphs is bounded by a polynomial of $|S|$ and the distance d . Namely, we prove the following theorem:

Theorem 26. *Consider any planar graph G , fix a non-empty set A containing $c \geq 1$ vertices of G . Then, the set $\{\pi_d[v, A] \mid v \in V(G)\}$ of all distance- d profiles on A has at most $64c^3(d+2)^7$ distinct elements.*

We remark that it is already known that in each class \mathcal{C} of graphs with bounded expansion the set $\{\pi_d[v, A] \mid v \in V(G)\}$ contains at most $f(d) \cdot |A|$ distinct elements for every non-empty set A , for some function f depending only on \mathcal{C} [38]. However, the proof of Reidl et al. requires f to grow exponentially, even for planar graphs. The polynomial dependence on d asserted by the Theorem 26 is crucial in the following parts of the proof.

The proof of Theorem 26 consists of two major parts. Firstly, we prove the restricted version of the theorem, called the Noose Profile Lemma (Theorem 28). Namely, we say that a non-empty subset A of vertices of a graph G is *nice* if there exists an embedding of G in the plane in which we can draw a simple closed curve (called *noose*) passing through each vertex of A exactly once, and not touching the interiors of edges at any point. The Noose Profile Lemma states that the statement of Theorem 26 holds for *nice* sets of vertices, with a slightly better polynomial dependence on $|A|$ and d .

Then, in Subsection 4.7, we lift the lemma to Theorem 26. The proof of this part has been inspired by the concept of “cutting open” the plane along the Steiner tree of a graph, present in the work of Pilipczuk et al. [32].

The theorem is then used as a black-box in the upper bound proof in Section 5, so Sections 4 and 5 can be read independently.

4.1 Statement of the Noose Profile Lemma

In order to state Noose Profile Lemma, we first need to define nooses in plane graphs.

Definition 27. For any graph G embedded in the plane, a **noose** is a closed curve in the plane that passes through vertices and faces of this graph that does not intersect itself or the interiors of any edges. Note that this definition disallows the curve to pass through any vertex multiple times.

We let the **length** of a noose to be the number of vertices on it. For any noose \mathcal{L} in a plane graph G , we define the set of vertices on the noose $V(\mathcal{L})$. Moreover, for each $v \in V(G)$, we define the distance- d profile of v on \mathcal{L} in G : $\pi_d[v, \mathcal{L}] := \pi_d[v, V(\mathcal{L})]$. Also, we define the subgraph $G_{\mathcal{L}}$ as the graph containing all vertices and edges within the bounded closed region whose boundary is \mathcal{L} (Figure 11). The vertices and edges of $G_{\mathcal{L}}$ are said to be **enclosed** by \mathcal{L} .

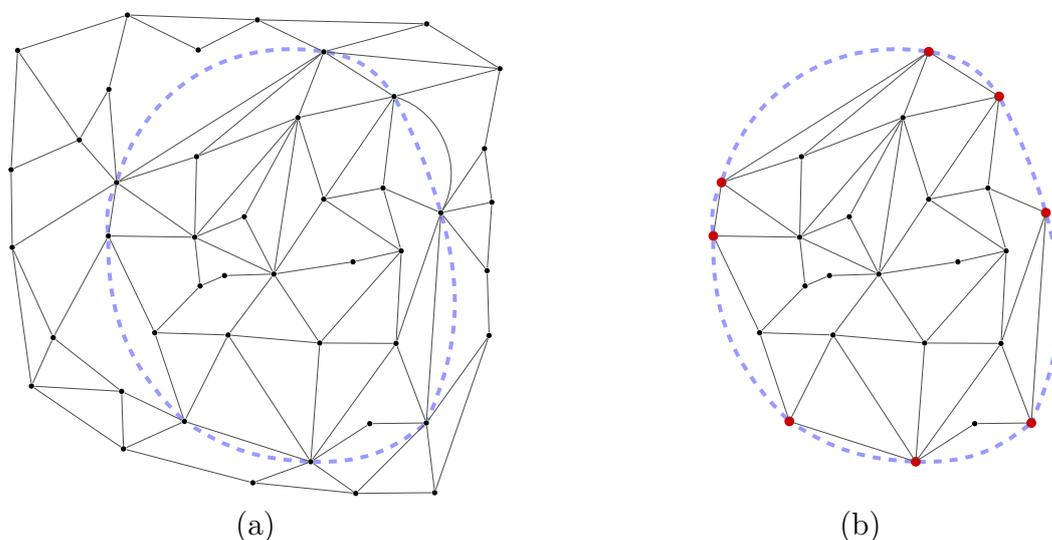


Figure 11: (a) — an example graph G with a noose \mathcal{L} (blue); (b) — the graph $G_{\mathcal{L}}$, with $V(\mathcal{L})$ marked red.

We are now ready to present the lemma in question:

Theorem 28 (Noose Profile Lemma). *Consider any graph G together with its embedding in the plane, and fix any noose \mathcal{L} containing $c \geq 1$ vertices of G . The set $\{\pi_d[v, \mathcal{L}] \mid v \in V(G_{\mathcal{L}})\}$ of different distance- d profiles on \mathcal{L} , measured from the vertices enclosed by \mathcal{L} , has at most $c^3(d+2)^4$ elements.*

We remark that if c is bounded by a polynomial of d , then the theorem above implies that the number of distance- d profiles on \mathcal{L} is also polynomially bounded on d .

Subsections 4.2 through 4.6 are devoted to the proof of the lemma.

4.2 Definition and properties of closeness

Let us first describe some helpful notation. For a closed region $\Omega \subseteq \mathbb{R}^2$ of the plane, we denote $\partial\Omega$ as the topological boundary of Ω , and $\text{Int } \Omega$ as its interior. We also introduce the following majorization relation on distance- d profiles in G :

Definition 29. For a graph G , a subset of vertices $S \subseteq V(G)$, and two vertices $s, t \in V(G)$, we say that the distance- d profile $\pi_d[s, S]$ **majorizes** $\pi_d[t, S]$ if $\pi_d[s, S](v) \geq \pi_d[t, S](v)$ for all $v \in S$. We denote this partial order as $\pi_d[s, S] \succcurlyeq \pi_d[t, S]$.

As in the statement of the lemma, fix any graph G and a noose \mathcal{L} . Let v_1, v_2, \dots, v_c be the vertices of \mathcal{L} , listed in the clockwise order. For simplicity, we set $v_0 = v_c$ and $v_{c+1} = v_1$. Let also $\mathcal{L}[i, j] \subseteq \mathcal{L}$ be the curve from v_i to v_j following the noose clockwise. That is, $\mathcal{L}[i, j]$ contains vertices v_i, v_{i+1}, \dots, v_j in this order if $i \leq j$, and $v_i, v_{i+1}, \dots, v_c, v_1, v_2, \dots, v_j$ otherwise. Moreover, let $\mathcal{L}(i, j) := \mathcal{L}[i, j] - \{v_i, v_j\}$.

Let us say that a vertex $s \in V(G_{\mathcal{L}})$ is **close** to a vertex $v_i \in \mathcal{L}$ if the distance between these two vertices in G does not exceed d , and no shortest path connecting v_i and s contains any vertex from \mathcal{L} as an internal vertex. Otherwise, we say that s is **far** from v_i . For such vertex s we also set $\text{Close}(s) = \{v \in \mathcal{L} \mid s \text{ is close to } v\}$.

Lemma 30. *If for any two vertices $s, t \in V(G_{\mathcal{L}})$ we have $\pi_d[s, \text{Close}(s)] \succcurlyeq \pi_d[t, \text{Close}(s)]$, then $\pi_d[s, \mathcal{L}] \succcurlyeq \pi_d[t, \mathcal{L}]$.*

Proof. Take any $v \in \mathcal{L}$. If also $v \in \text{Close}(s)$, then

$$\pi_d[s, \mathcal{L}](v) = \pi_d[s, \text{Close}(s)](v) \geq \pi_d[t, \text{Close}(s)](v) = \pi_d[t, \mathcal{L}](v).$$

If however $v \notin \text{Close}(s)$, then either of the following cases holds:

- The distance between s and v is longer than d . In this case, $\pi_d[s, \mathcal{L}](v) = +\infty \geq \pi_d[t, \mathcal{L}](v)$.
- The shortest path between s and v has length at most d , and some such shortest path passes through another vertex of \mathcal{L} . Pick a vertex of the noose $v' \in \mathcal{L}$ on any such shortest path. In case there are multiple such vertices, we pick any vertex minimizing its distance from s . As no shortest path between s and v' contains any other vertex of \mathcal{L} (otherwise, such vertex would be even closer to s), we have that $v' \in \text{Close}(s)$. Hence,

$$\begin{aligned} \pi_d[s, \mathcal{L}](v) &= \text{dist}(s, v) = \text{dist}(s, v') + \text{dist}(v', v) = \pi_d[s, \text{Close}(s)](v') + \text{dist}(v', v) \geq \\ &\geq \pi_d[t, \text{Close}(s)](v') + \text{dist}(v', v) = \text{dist}(t, v') + \text{dist}(v', v) \geq \\ &\geq \text{dist}(t, v) = \pi_d[t, \mathcal{L}](v). \end{aligned}$$

In all cases, we proved that $\pi_d[s, \mathcal{L}](v) \geq \pi_d[t, \mathcal{L}](v)$. Thus, $\pi_d[s, \mathcal{L}] \succcurlyeq \pi_d[t, \mathcal{L}]$. □

This immediately leads to the following corollary:

Corollary 31. *If for any two vertices $s, t \in V(G_{\mathcal{L}})$ we have*

$$\text{Close}(s) = \text{Close}(t) \quad \text{and} \quad \pi_d[s, \text{Close}(s)] = \pi_d[t, \text{Close}(t)],$$

then $\pi_d[s, \mathcal{L}] = \pi_d[t, \mathcal{L}]$.

Hence, $\pi_d[s, \mathcal{L}]$ can be uniquely deduced from the distances between s and each vertex from $\text{Close}(s)$. This allows us to safely replace the graph G with $G_{\mathcal{L}}$ (that is, to remove all vertices and edges not enclosed by \mathcal{L}); this replacement does not change $\text{Close}(s)$ or any of the distances to the close vertices.

In $G_{\mathcal{L}}$, let $\mathcal{C}_i := \{s \in V(G_{\mathcal{L}}) \mid |\text{Close}(s)| = i\}$. First, we examine the behavior of the vertices from \mathcal{C}_0 and \mathcal{C}_1 . Note that $\mathcal{L} \subseteq \mathcal{C}_1$.

Lemma 32. *There are at most $c(d+1) + 1$ distinct distance- d profiles on \mathcal{L} measured from the vertices in $\mathcal{C}_0 \cup \mathcal{C}_1$.*

Proof. Obviously, no vertex from \mathcal{C}_0 can reach \mathcal{L} by a path consisting of at most d edges, so each such vertex generates the same distance- d profile on \mathcal{L} : the constant function $\lambda v.(+\infty)$.

Now consider any vertex $s \in \mathcal{C}_1$. Its distance- d profile on \mathcal{L} can be unambiguously deduced from the only vertex $v_i \in \text{Close}(s)$ and the distance $\text{dist}(s, v_i) \in [0, d]$ (Corollary 31). This leads to a maximum of $c(d+1)$ distinct profiles created by \mathcal{C}_1 .

Hence, at most $c(d+1) + 1$ distinct distance- d profiles can be generated by the vertices in $\mathcal{C}_0 \cup \mathcal{C}_1$. \square

From now on, we only consider remaining vertices.

4.3 Rightmost shortest paths

We will now extend slightly the definitions by Klein [16].

Definition 33. Fix any vertex $v \in \mathcal{L}$. We define a strict partial order \ll on paths originating from v , but not necessarily terminating at the same vertex. We say that $P_1 \ll P_2$ (or: “ **P_1 is left of P_2** ”) if:

- neither of the paths is a prefix of another,
- P_1 and P_2 have a common prefix x_1, x_2, \dots, x_k , where $v = x_1$,
- w_i immediately succeeds x_k in P_i for $i = 1, 2$,
- edge $x_k w_1$ is “to the left” of edge $x_k w_2$; formally, edges $x_k x_{k-1}$, $x_k w_1$ and $x_k w_2$ are in the clockwise order around x_k (Figure 12(a)). However, if $k = 1$, the paths do not contain x_0 . In order to cope with this problem, we imagine an auxiliary vertex x_0 outside of the region enclosed by \mathcal{L} and connect it by an edge to x_1 , placed strictly outside of this region (Figure 12(b)).

We can also interchangeably say $P_2 \gg P_1$ (or: “ **P_2 is right of P_1** ”).

The relation \ll was defined in [16] only between the simple paths originating at v and terminating at a fixed vertex s . Then \ll is a linear order. In that setup, if v and s are in the same connected component of the graph, this relation limited only to the shortest paths between v and s contains the maximal element $P_{v,s}$, which we call the **rightmost shortest path**. However, Definition 33 naturally extends \ll to any two paths P_1, P_2 which both originate at v and neither is a prefix of the other, but which do not have to terminate at the same vertex.

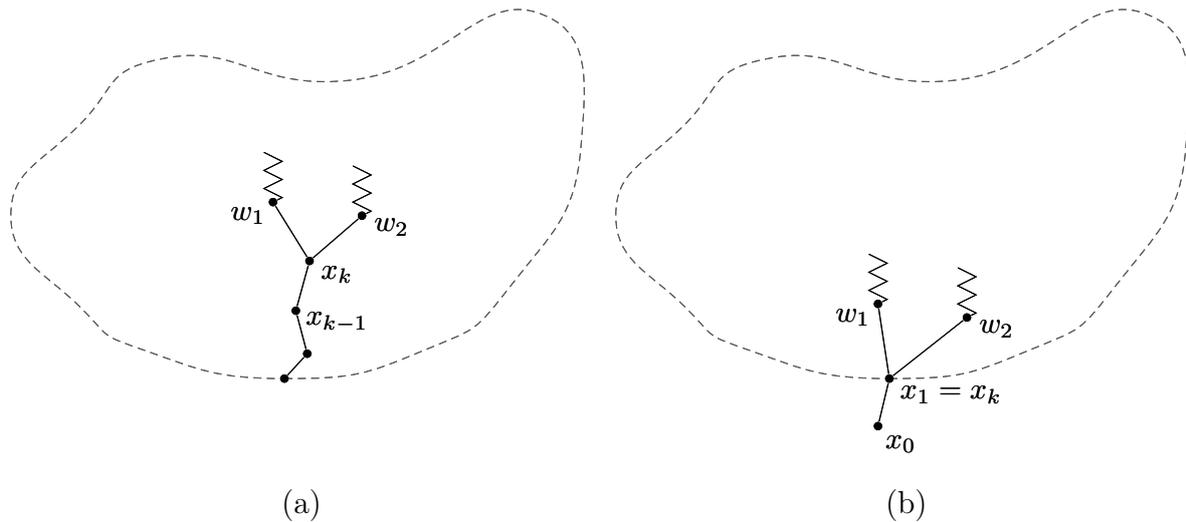


Figure 12: (a) — the path containing w_1 is to the left of the path containing w_2 .
 (b) — if the first edges of two compared paths differ, we imagine an auxiliary vertex x_0 and an auxiliary edge x_0x_1 which we prepend to both paths. We then compare the paths as in (a).

We also define the rightmost shortest path tree rooted at v as in [16]. Notably, for any vertex s of the tree, the path between v and s in this tree is $P_{v,s}$. Obviously, if $v \in \mathcal{L}$ and for some vertices $a, b \in V(G_{\mathcal{L}})$, a belongs to $P_{v,b}$, then $P_{v,a} = P_{v,b}[v, a]$. Moreover, for $v \in \mathcal{L}$ and $s \in V(G_{\mathcal{L}})$, we obviously have that $\text{len}(P_{v,s}) = \text{dist}(v, s)$.

Lemma 34. *For any two distinct vertices $s, t \in \mathcal{L}$, paths $P_{a,s}$ and $P_{a,t}$ share a common prefix and are vertex-disjoint after removing this prefix.*

Proof. Consider the rightmost shortest path tree rooted at a . The tree contains paths $P_{a,s}$ and $P_{a,t}$. The root a belongs to both paths, hence they both must intersect at a common prefix and nowhere else. \square

Lemma 35. *Fix any $a, b \in \mathcal{L}$ and $s \in V(G_{\mathcal{L}}) - V(\mathcal{L})$ such that $a, b \in \text{Close}(s)$. Paths $P_{a,s}$ and $P_{b,s}$ share a common suffix as a subpath, and are vertex-disjoint after removing this suffix.*

Proof. Take the first vertex $x \in P_{a,s} \cap P_{b,s}$ on the path between a and s , that is, the one minimizing $\text{dist}(x, a)$. If $x = s$, we are done. In the opposite case, we set a^* and b^* as the next vertices after x on the paths $P_{a,s}$ and $P_{b,s}$, respectively.

Moreover, we set a_p and b_p as the last vertices before x on these paths. Vertices a_p and b_p exist since $P_{a,s}$ cannot pass through b (otherwise, s would not be close to a), and similarly $P_{b,s}$ cannot pass through a . Hence, $x \notin \{a, b\}$ — that is, x cannot be the initial vertex of either of these paths.

Obviously, the length of each of the paths $P_{a,s}[x, s]$ and $P_{b,s}[x, s]$ must be the same, otherwise we would be able to swap one of the suffixes with another in one of the paths $P_{a,s}, P_{b,s}$, reducing its length. This is, however, impossible as we require $P_{a,s}$ and $P_{b,s}$ to be the shortest paths connecting the corresponding pairs of vertices. Also, note that paths $P_{a,x}$ and $P_{b,x}$ are vertex-disjoint apart from their common endpoint x .

Consider the closed region Ω bounded by $P_{a,x}$, $P_{b,x}$ and the part of the noose between a and b — either $\mathcal{L}[a, b]$ or $\mathcal{L}[b, a]$ — chosen so that $s \in \text{Int } \Omega$. (Obviously, s cannot belong to \mathcal{L} , $P_{a,x}$ or $P_{b,x}$.) Assume now that $a^* \notin \Omega$ (Figure 13(a)). Then, the path $P_{a,s}[a^*, s]$ must intersect $\partial\Omega$. Hence, one of the following cases occurs:

- The path intersects $P_{a,x}$ (Figure 13(b)). In this case, the path $P_{a,s}$ would intersect itself.
- The path intersects $P_{b,x}$ at some point $y \neq x$ (Figure 13(c)). In this case, $P_{a,s}$ contains vertices x, y in this order, while $P_{b,s}$ contains these vertices in the opposite order. Hence,

$$\text{len}(P_{a,s}[y, s]) < \text{len}(P_{a,s}[x, s]) = \text{len}(P_{b,s}[x, s]) < \text{len}(P_{b,s}[y, s]).$$

Therefore, $P_{b,s}$ could be shortened by replacing its suffix $P_{b,s}[y, s]$ by $P_{a,s}[y, s]$.

- The path enters Ω through an edge whose either endpoint is in $\mathcal{L} - \{a, b\}$. This is impossible as this edge would not be enclosed by the noose and therefore would not belong to $G_{\mathcal{L}}$.

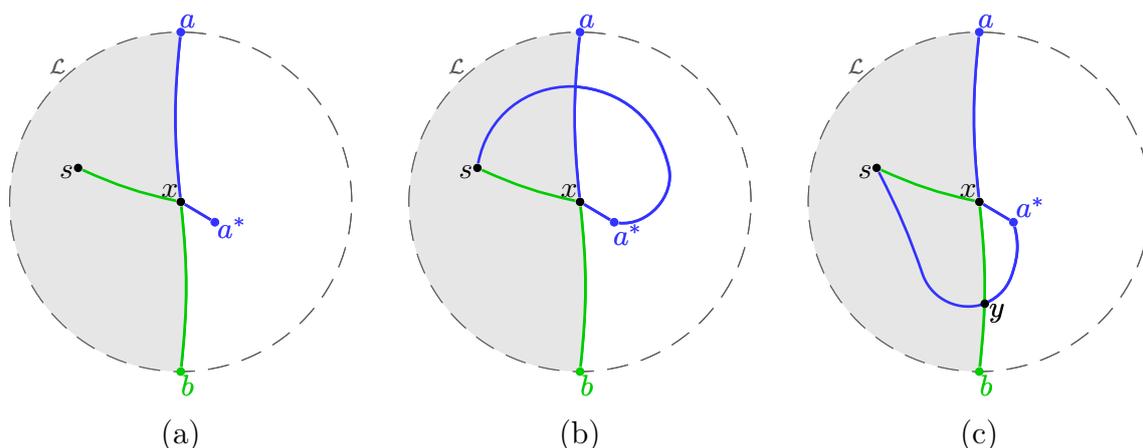


Figure 13: (a) — an example setup. The noose \mathcal{L} is a gray dashed circle. The region Ω is shaded. Green path is $P_{b,s}$, and the blue path is the prefix of $P_{a,s}$.

(b) — a case where $P_{a,s}[a^*, s]$ intersects $P_{a,x}$.

(c) — a case where $P_{a,s}[a^*, s]$ intersects $P_{b,x}$.

We arrived at a contradiction in each of the cases above. Therefore $a^* \in \Omega$, and analogously $b^* \in \Omega$.

Assume now that $a^* \neq b^*$. Then, we remind that $P_{a,s}$ is rightmost among all shortest paths from a to s . Consider the path $Q_a = P_{a,x} \cdot P_{b,s}[x, s]$. As $\text{len}(P_{a,s}[x, s]) = \text{len}(P_{b,s}[x, s])$, we also have that $\text{len}(Q_a) = \text{len}(P_{a,s})$. Since $Q_a \neq P_{a,s}$, we get that $Q_a \ll P_{a,s}$. Notice that a_p and x are two last common vertices of Q_a and $P_{a,s}$, while b^* is the first vertex of Q_a not belonging to $P_{a,s}$, and a^* is the first vertex of $P_{a,s}$ not belonging to Q_a . By the definition of \ll , edges xa_p , xb^* and xa^* form a clockwise order around x . We similarly deduce that edges xb_p , xa^* and xb^* form a clockwise order around x . This implies the clockwise order of four edges around x : xa_p, xb^*, xb_p, xa^* . As $a_p x$ and xb_p are

two consecutive edges on the boundary of Ω , we conclude that this boundary separates a^* and b^* . That is however impossible as $a^*, b^* \in \Omega$.

From the contradiction above, $a^* = b^*$. This means that $P_{a,s}$ and $P_{b,s}$ share an edge xa^* . We now prove that they must share the remaining suffix of the path from x to s . Otherwise, there would be an edge fg shared by both $P_{a,s}$ and $P_{b,s}$, followed by an edge gh_a in $P_{a,s}$, and by gh_b in $P_{b,s}$ ($h_a \neq h_b$). Assume without loss of generality that the edges fg, gh_a, gh_b are oriented clockwise around g . By the Definition 33, this means that

$$P_{a,s} = P_{a,s}[a, g] \cdot P_{a,s}[g, s] \ll P_{a,s}[a, g] \cdot P_{b,s}[g, s]$$

— a contradiction. Hence, these paths must share the remaining suffix of the path from x to s .

As $P_{a,s}[a, x] \cap P_{b,s}[b, x] = \{x\}$ and $P_{a,s}[x, s] = P_{b,s}[x, s]$, this finishes the proof. \square

Fix any vertex $s \in \mathcal{C}_k$, $k \geq 2$ and assume that it is close to vertices $v_{i_1}, v_{i_2}, \dots, v_{i_k}$ of the noose, listed in the clockwise order ($1 \leq i_1 < i_2 < \dots < i_k \leq c$). For simplicity, we also denote $i_{k+1} = i_1$. Now, consider k paths $P_{v_{i_1},s}, P_{v_{i_2},s}, \dots, P_{v_{i_k},s}$.

Lemma 36. *Exactly one of the following conditions holds:*

- (1) *There exists $j \in [1, k]$ such that paths $P_{v_{i_j},s}$ and $P_{v_{i_{j+1}},s}$ are vertex disjoint apart from s ,*
- (2) *There exists a vertex $s' \neq s$ such that all the paths $P_{v_{i_1},s}, P_{v_{i_2},s}, \dots, P_{v_{i_k},s}$ contain s' . This vertex is adjacent to s and is the penultimate vertex in each of these paths.*

Proof. Obviously, condition (2) implies that (1) is false. We only need to prove the opposite implication.

Assume that property (1) holds for no $j \in [1, k]$. Therefore, by Lemma 35, we see that for each $j \in [1, k]$, paths $P_{v_{i_j},s}$ and $P_{v_{i_{j+1}},s}$ share the common suffix, containing at least two vertices. Hence, each path $P_{v_{i_1},s}, P_{v_{i_2},s}, \dots, P_{v_{i_k},s}$ shares the same common suffix, having at least two vertices. The penultimate vertex in this suffix is s' as in property (2). \square

Let V_1 be the set of vertices from $\mathcal{C}_2 \cup \mathcal{C}_3 \cup \dots \cup \mathcal{C}_c$ for which the first condition in Lemma 36 holds, and V_2 — the analogous set for the second condition.

Lemma 37. *For every $s \in V_2$, there exists a vertex $s' \in V_1$ and an integer $m \in \{1, 2, \dots, d\}$ such that*

$$\forall_{v \in \mathcal{L}} \pi_d[s, \mathcal{L}](v) = \text{Cap}_d(\pi_d[s', \mathcal{L}](v) + m),$$

where

$$\text{Cap}_d(x) = \begin{cases} x & \text{if } x \leq d, \\ +\infty & \text{otherwise.} \end{cases}$$

Proof. If $s \in V_2$, then there exists a sequence $(s_0, s_1, \dots, s_{m-1}, s_m)$ where $s = s_m$ and $m \geq 1$, which is the longest common suffix of all the paths $P_{v_{i_1},s}, P_{v_{i_2},s}, \dots, P_{v_{i_k},s}$. The rightmost shortest paths from $v_{i_1}, v_{i_2}, \dots, v_{i_k}$ to s_0 can be obtained by dropping m last vertices in each of the paths. In this setup, we have $s_0 \in V_1$ — otherwise, we would

be able to prepend the sequence defined above by the common neighbor of s_0 on all the paths terminating at s_0 . We will now prove that $s' = s_0$ and m satisfy the condition stated in the lemma.

Pick any vertex $v \in \mathcal{L}$ and consider the following cases:

- If s is close to v , then $\text{dist}(s, v) = \text{dist}(s_0, v) + m$ (because of the common prefix above).
- If s is far from v and $\pi_d[s, \mathcal{L}](v) = \infty$, then $\text{dist}(s, v) > d$. Hence, $\text{dist}(s_0, v) > d - m$ as $\text{dist}(s, v) \leq \text{dist}(s, s_0) + \text{dist}(s_0, v) = m + \text{dist}(s_0, v)$. To put it in the other way, $\text{dist}(s_0, v) + m > d$. Therefore $\pi_d[s_0, \mathcal{L}](v) + m > d$, which is consistent with the condition $\pi_d[s, \mathcal{L}](v) = \infty$ required by the statement of the lemma.
- If s is far from v and $\pi_d[s, \mathcal{L}](v) \leq d$, then by definition of closeness some shortest path from s to v passes through another vertex $v' \in V(\mathcal{L}) - \{v\}$. Assume that v' is one of the vertices closest to s with this property.

We see that s is close to v' . Thus, $v' \in \{v_{i_1}, \dots, v_{i_k}\}$, so $P_{v', s}$ is one of the paths in the set $\{P_{v_{i_1}, s}, \dots, P_{v_{i_k}, s}\}$, and as such, it contains s_0 . Therefore, some shortest path from v to s contains v, v', s_0, s in this order. Hence, $\pi_d[s, \mathcal{L}](v) = \text{dist}(s, v) = \text{dist}(s, s_0) + \text{dist}(s_0, v) = m + \text{dist}(s_0, v) = \pi_d[s_0, \mathcal{L}](v) + m$.

Therefore, the distance- d profile $\pi_d[s, \mathcal{L}]$ is constructed by considering the profile $\pi_d[s_0, \mathcal{L}]$ where $s_0 \in V_1$, and increasing each entry in the profile by m , while remembering to replace each integer exceeding d by ∞ . □

Corollary 38. *If*

$$|\{\pi_d[v, \mathcal{L}] \mid v \in V_1\}| = B \quad \text{for an integer } B,$$

then

$$|\{\pi_d[v, \mathcal{L}] \mid v \in V_1 \cup V_2\}| \leq B \cdot (d + 1).$$

Proof. By Lemma 37, the profile of a vertex $v \in V_2$ can be deduced from a distance- d profile of a vertex $v' \in V_1$ and an integer $m \in \{1, 2, \dots, d\}$. Hence,

$$|\{\pi_d[v, \mathcal{L}] \mid v \in V_2\}| \leq |\{\pi_d[v, \mathcal{L}] \mid v \in V_1\}| \cdot |\{1, 2, \dots, d\}| = B \cdot d.$$

We also include the vertices in V_1 into this inequality and get:

$$\begin{aligned} |\{\pi_d[v, \mathcal{L}] \mid v \in V_1 \cup V_2\}| &\leq |\{\pi_d[v, \mathcal{L}] \mid v \in V_1\}| + |\{\pi_d[v, \mathcal{L}] \mid v \in V_2\}| \leq \\ &\leq B \cdot (d + 1). \end{aligned} \quad \square$$

Corollary 38 allows us to focus on bounding the number of different distance- d profiles measured only from the vertices in V_1 .

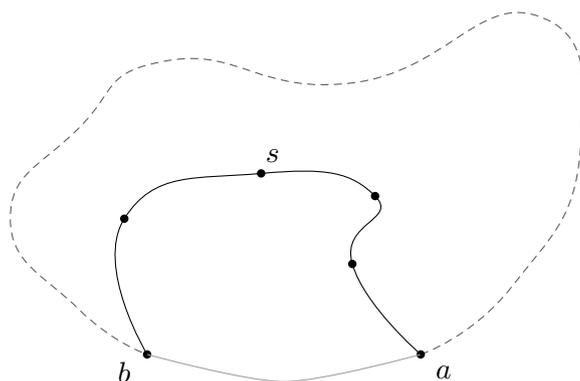


Figure 14: Definition of a bucket assigned to a vertex s . Here, $d_a = 3$ and $d_b = 2$. Vertex s is far from any vertex on the dashed part of the noose.

4.4 Buckets

We will now group all the vertices in V_1 into a reasonably small number of buckets. The buckets will be chosen in such a way that all the vertices within one group expose similar properties in G . This will eventually allow us to bound the number of different distance- d profiles measured by each individual bucket.

Definition 39. For a vertex $s \in V_1$, we define the **bucket** of s as any quadruple (a, b, d_a, d_b) satisfying the following conditions (Figure 14):

- $a, b \in V(\mathcal{L})$ and $a, b \in \text{Close}(s)$,
- $\mathcal{L}(b, a) \cap \text{Close}(s) = \emptyset$; in other words, each vertex $v \in \mathcal{L}(b, a)$ is far from s ,
- paths $P_{a,s}$ and $P_{b,s}$ intersect only at s ,
- $d_a = \text{dist}(a, s)$ is the distance between a and s in $G_{\mathcal{L}}$, and $1 \leq d_a \leq d$,
- $d_b = \text{dist}(b, s)$ is the distance between b and s in $G_{\mathcal{L}}$, and $1 \leq d_b \leq d$.

Lemma 36 proves that such an assignment exists since $s \in V_1$.

Notice that there at most $c(c-1)d^2$ possible quadruples (a, b, d_a, d_b) — a and b are two distinct vertices of the noose in some order, and d_a, d_b are two positive distances not exceeding d .

Fix a bucket (a, b, d_a, d_b) and consider the set $\mathcal{S} = \mathcal{S}(a, b, d_a, d_b)$ of vertices s that are assigned to the bucket. Remember that each $s \in \mathcal{S}$ is far from each vertex of $\mathcal{L}(b, a)$.

Lemma 40. For any two different vertices $s, t \in \mathcal{S}$, if paths $P_{a,t}$ and $P_{b,s}$ intersect at any vertex x , then $P_{a,s} \gg P_{a,t}$, $P_{b,t} \gg P_{b,s}$, and $\text{dist}(x, t) = \text{dist}(x, s) > 0$ (Figure 15).

Proof. Assume that $P_{a,t}$ and $P_{b,s}$ intersect at x . Then we have a couple of simple equalities:

$$\text{dist}(a, t) = \text{dist}(a, x) + \text{dist}(x, t) \quad \text{and} \quad \text{dist}(b, s) = \text{dist}(b, x) + \text{dist}(x, s).$$

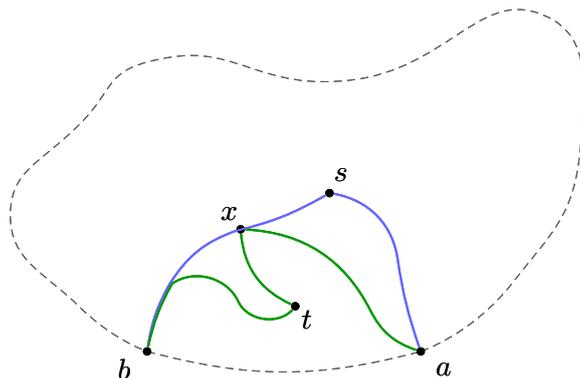


Figure 15: A possible configuration in Lemma 40. Paths $P_{a,s}$ and $P_{b,s}$ are marked blue, while $P_{a,t}$ and $P_{b,t}$ are green.

From the triangle inequality and the fact that $s, t \in \mathcal{S}(a, b, d_a, d_b)$, we infer:

$$\text{dist}(a, x) + \text{dist}(x, s) \geq \text{dist}(a, s) = d_a = \text{dist}(a, t) = \text{dist}(a, x) + \text{dist}(x, t), \quad (2)$$

$$\text{dist}(b, x) + \text{dist}(x, s) = \text{dist}(b, s) = d_b = \text{dist}(b, t) \leq \text{dist}(b, x) + \text{dist}(x, t). \quad (3)$$

We infer that $\text{dist}(x, s) \geq \text{dist}(x, t)$ from (2) and $\text{dist}(x, s) \leq \text{dist}(x, t)$ from (3). Therefore, $\text{dist}(x, s) = \text{dist}(x, t)$ and all the inequalities above are satisfied with equality. As $x \neq s$ (otherwise we would have $x = s = t$), obviously $\text{dist}(x, s) > 0$.

Hence, $\text{dist}(a, s) = \text{dist}(a, x) + \text{dist}(x, s)$ and $\text{dist}(b, t) = \text{dist}(b, x) + \text{dist}(x, t)$.

We now see two shortest paths from a to s : the rightmost path $P_{a,s}$, and another path $Q := P_{a,x} \cdot P_{b,s}[x, s]$.

Notice that $P_{a,s}$ cannot be the prefix of $P_{a,x}$ as it is strictly longer than $P_{a,x}$. Also $P_{a,x}$ cannot be the prefix of $P_{a,s}$ because x lies on $P_{b,s}$ (which is vertex disjoint with $P_{a,s}$ apart from s , and $x \neq s$). Hence, we have either $P_{a,s} \ll P_{a,x}$ or $P_{a,s} \gg P_{a,x}$. We however cannot have $P_{a,s} \ll P_{a,x}$, or otherwise we would have $P_{a,s} \ll P_{a,x} \cdot P_{b,s}[x, s] = Q$, which is impossible by the definition of $P_{a,s}$. Therefore $P_{a,s} \gg P_{a,x}$. As $P_{a,x} = P_{a,t}[a, x]$, we also get $P_{a,s} \gg P_{a,t}$.

An analogous argument shows that $P_{b,t} \gg P_{b,s}$. □

Corollary 41. *For any two distinct vertices $s, t \in \mathcal{S}$, if paths $P_{a,s}$ and $P_{b,t}$ intersect at any vertex x , then $P_{a,t} \gg P_{a,s}$, $P_{b,s} \gg P_{b,t}$, and $\text{dist}(x, t) = \text{dist}(x, s) > 0$.*

Proof. In the statement of Lemma 40, swap s and t . □

Corollary 42. *It is impossible for $P_{a,t}$ and $P_{b,s}$ to intersect, and for $P_{a,s}$ and $P_{b,t}$ to intersect at the same time.*

Proof. The former condition would imply $P_{a,s} \gg P_{a,t}$ (Lemma 40), while the latter — $P_{a,t} \gg P_{a,s}$ (Corollary 41). □

4.5 Gamma-regions

We fix any bucket $\mathcal{S} = \mathcal{S}(a, b, d_a, d_b)$. We will now assign a region of the plane to each vertex $w \in \mathcal{S}$.

Lemma 43. *For any vertex $w \in \mathcal{S}$, the curve $\mathcal{L}[a, b] \cup P_{a,w} \cup P_{b,w}$ is a closed Jordan curve.*

Proof. Notice that $w \notin V(\mathcal{L})$ and w is close to a , so $P_{a,w}$ cannot contain vertices from \mathcal{L} as internal vertices. Therefore, $P_{a,w} \cap \mathcal{L}[a, b] = \{a\}$. Analogously, $P_{b,w} \cap \mathcal{L}[a, b] = \{b\}$. Finally, $P_{a,w} \cap P_{b,w} = \{w\}$ due to the disjointness of these paths in the sense of Lemma 36. Thus the union of these curves is a closed Jordan curve. \square

Definition 44. For any vertex $w \in \mathcal{S}$, we define the **gamma-region** Γ_w as the closed, bounded region of the plane whose boundary is $\mathcal{L}[a, b] \cup P_{a,w} \cup P_{b,w}$ (Figure 16).

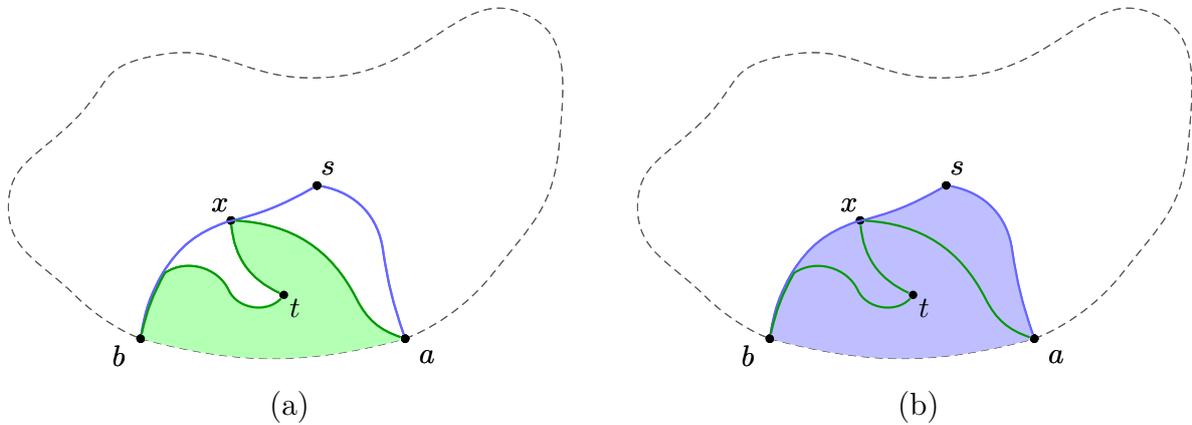


Figure 16: Gamma regions in the configuration from Figure 15: (a) — Γ_t , (b) — Γ_s .

We will now prove the potential of this definition. Intuitively, for any two vertices $s, t \in \mathcal{S}$ in the bucket, we will either have $t \in \Gamma_s$, which will imply a serious structural relation between the gamma-regions Γ_s and Γ_t , or we will have $t \notin \Gamma_s$, which will in turn have major implications for the distance- d profiles $\pi_d[s, \mathcal{L}]$ and $\pi_d[t, \mathcal{L}]$.

Lemma 45. *For any two different vertices $s, t \in \mathcal{S}$, if $t \in \Gamma_s$, then $\Gamma_t \subsetneq \Gamma_s$.*

Proof. Assume that $t \in \Gamma_s$. We will first prove a couple of helper claims.

Claim 46. *If $P_{b,t}$ is disjoint with $P_{a,s}$, then $P_{b,t} \gg P_{b,s}$ and $P_{b,t}$ is fully contained within Γ_s .*

Proof. Assume that $P_{b,t}$ is disjoint with $P_{a,s}$. As $\text{dist}(b, t) = \text{dist}(b, s)$ and $s \neq t$, neither of paths $P_{b,s}$ and $P_{b,t}$ is a prefix of another and thus they are ordered by \ll . We notice that $P_{b,s}$ follows $\partial\Gamma_s$ clockwise. This is because we defined $\mathcal{L}[a, b]$ to run clockwise from a to b on \mathcal{L} , and $P_{b,s}$ is a directed path originating from b , where $\mathcal{L}[a, b]$ terminated. Hence, both $P_{b,s}$ and $\mathcal{L}[a, b]$ have to be oriented the same way on $\partial\Gamma_s$.

Let $(y_1, y_2, \dots, y_{k-1}, y_k)$ ($y_k = t$) be the maximal suffix of $P_{b,t}$ which is disjoint with $P_{b,s}$. For any choice of $i \in \{1, 2, \dots, k\}$, we have that $y_i \notin P_{a,s}$ since $P_{b,t}$ is disjoint with

$P_{a,s}$, and $y_i \notin \mathcal{L}[a, b]$ as $\mathcal{L}[a, b] \cap P_{b,t} = \{b\}$. Hence $\{y_1, y_2, \dots, y_k\} \cap \partial\Gamma_s = \emptyset$, and thus these vertices are either all strictly inside of Γ_s , or all strictly outside of it.

Now, as $P_{b,s}$ follows the boundary of Γ_s clockwise, it can be easily seen that $y_1 \in \Gamma_s \Leftrightarrow P_{b,y_1} \gg P_{b,s}$; in other words, the interior of Γ_s is “to the right” of $P_{b,s}$.

If $P_{b,y_1} \gg P_{b,s}$, then $P_{b,t} \gg P_{b,s}$ and $P_{b,s}$ is fully within Γ_s .

If $P_{b,y_1} \ll P_{b,s}$, then $P_{b,t} \ll P_{b,s}$. Because $y_1 \notin \Gamma_s$, we have that $\{y_1, y_2, \dots, y_k\}$ is disjoint with Γ_s . In particular, $t \notin \Gamma_s$. This however contradicts our assumption that $t \in \Gamma_s$. \lrcorner

Claim 47. *If $P_{a,t}$ is disjoint with $P_{b,s}$, then $P_{a,t} \ll P_{a,s}$ and $P_{a,t}$ is fully contained within Γ_s .*

Proof. The proof of this claim follows mostly the proof of Claim 46. However, $P_{a,s}$ follows $\partial\Gamma_s$ anti-clockwise. This is because $\mathcal{L}[a, b]$ runs clockwise along the noose from a to b , and both $\mathcal{L}[a, b]$ and $P_{a,s}$ originate from the same vertex a . Hence, these two paths are oriented in the opposite way along $\partial\Gamma_s$. This means that the interior of Γ_s is “to the left” of $P_{a,s}$. Hence, $t \in \Gamma_s$ if and only if $P_{a,t} \ll P_{a,s}$. \lrcorner

Note that if both $P_{b,t}$ is disjoint with $P_{a,s}$, and $P_{a,t}$ is disjoint with $P_{b,s}$, then by Claims 46 and 47 we get that $\partial\Gamma_t \subseteq \Gamma_s$ and thus $\Gamma_t \subseteq \Gamma_s$.

Assume that $t \in \Gamma_s$, but $\Gamma_t \not\subseteq \Gamma_s$. As $\partial\Gamma_t$ is a Jordan curve, there exists a vertex y on $\partial\Gamma_t$ not belonging to Γ_s . We will consider multiple cases, depending on which part of $\partial\Gamma_t$ this vertex belongs to.

Case 1: $y \in \mathcal{L}[a, b]$. This is however impossible as $\mathcal{L}[a, b] \subseteq \Gamma_s$.

Case 2a: $y \in P_{b,t}$ and $P_{b,t} \gg P_{b,s}$. In this case, $P_{b,t}$ cannot intersect $P_{a,s}$ by Corollary 41. Hence, by Claim 46, we have that $P_{b,t}$ is fully contained within Γ_s — a contradiction because $y \in P_{b,t}$ must lie outside of Γ_s .

Case 2b: $y \in P_{b,t}$ and $P_{b,t} \ll P_{b,s}$. In this case, Claim 46 requires that the intersection between $P_{b,t}$ and $P_{a,s}$ must exist. Hence, we also get $P_{a,t} \gg P_{a,s}$ by Corollary 41.

Now, $P_{a,t}$ cannot intersect $P_{b,s}$ as this is forbidden by Corollary 42. This leads to a contradiction due to Claim 47.

Case 3a: $y \in P_{a,t}$ and $P_{a,t} \gg P_{a,s}$. Similarly to the first part of Case 2b, we deduce that $P_{a,t}$ must intersect $P_{b,s}$. However, by Lemma 40, we get that $P_{a,t} \ll P_{a,s}$ — a contradiction.

Case 3b: $y \in P_{a,t}$ and $P_{a,t} \ll P_{a,s}$. For our convenience, we take y to be the first vertex of $P_{a,t}$ not in Γ_s . We again deduce that $P_{a,t}$ intersects $P_{b,s}$; otherwise, Claim 47 would imply that $P_{a,t}$ is contained within Γ_s , which contradicts the choice of y .

Let x_1 and x_2 be the first and the last intersection of $P_{a,t}$ with $P_{b,s}$, respectively. Notice that y cannot precede x_1 on $P_{a,t}$; as $P_{a,t} \ll P_{a,s}$, then $P_{a,y} \ll P_{a,s}$. As $P_{a,y}$ would not intersect $P_{b,s}$, we could prove (analogously to the proofs of Claims 46 and 47) that $P_{a,y}$ is fully contained within Γ_s , which contradicts our assumption. Vertex y cannot also appear after x_2 on $P_{a,t}$ — in this case, $y \notin \Gamma_s$ would imply $t \notin \Gamma_s$ (note that $P_{a,t}[y, t]$ would not be able to intersect $\partial\Gamma_s$), which again is impossible.

Since vertices a, x_1, y, x_2, t appear on $P_{a,t}$ in this order, we have that $x_1 \in P_{a,x_2}$. Also, by Lemma 40 and the fact that $\text{dist}(x_2, t) < \text{dist}(x_1, t)$, we get that

$$\begin{aligned} \text{dist}(b, x_1) &= \text{dist}(b, s) - \text{dist}(x_1, s) = \text{dist}(b, s) - \text{dist}(x_1, t) < \text{dist}(b, s) - \text{dist}(x_2, t) = \\ &= \text{dist}(b, s) - \text{dist}(x_2, s) = \text{dist}(b, x_2). \end{aligned}$$

We infer that $x_1 \in P_{b,s}[b, x_2] = P_{b,x_2}$.

Now, consider two paths: $P_{a,x_2} = P_{a,t}[a, x_2]$ and $P_{b,x_2} = P_{b,s}[b, x_2]$. By Lemma 35, these two paths must intersect on their common suffix and nowhere else. As $x_1 \in P_{a,x_2} \cap P_{b,x_2}$, we get that $P_{a,x_2}[x_1, x_2] = P_{b,x_2}[x_1, x_2]$. This means that $P_{b,s}$ shares its vertices with $P_{a,t}$ on the segment between x_1 and x_2 . Hence,

$$P_{a,t}[x_1, x_2] = P_{a,x_2}[x_1, x_2] = P_{b,x_2}[x_1, x_2] = P_{b,s}[x_1, x_2] \subseteq P_{b,s} \subseteq \partial\Gamma_s.$$

However, we get a contradiction because of our assumptions that $y \in P_{a,x_2}[x_1, x_2]$ and $y \notin \Gamma_s$.

After considering all cases above, we have that $\Gamma_t \subseteq \Gamma_s$. Moreover, $t \notin P_{a,s}$ (since $\text{dist}(a, s) = \text{dist}(a, t)$) and analogously $t \notin P_{b,s}$, and $t \notin V(\mathcal{L})$ (otherwise t would only be close to itself). Thus, $t \notin \partial\Gamma_s$. Since obviously $t \in \partial\Gamma_t$, we immediately get that $\Gamma_s \neq \Gamma_t$. Therefore, $\Gamma_t \subsetneq \Gamma_s$. \square

The claim above described the case $t \in \Gamma_s$. Now we try to see what happens when $t \notin \Gamma_s$.

Lemma 48. *If $t \notin \Gamma_s$ for some $s, t \in \mathcal{S}$, then $\pi_d[t, \mathcal{L}] \succcurlyeq \pi_d[s, \mathcal{L}]$.*

Proof. We will prove that $\pi_d[t, \text{Close}(t)] \succcurlyeq \pi_d[s, \text{Close}(t)]$. Lemma 30 will then conclude our proof.

Recall, by the choice of a and b , that $\text{Close}(t) \subseteq \mathcal{L}[a, b]$. We will now prove that $\text{dist}(v, s) \geq \text{dist}(v, t)$ for each $v \in \mathcal{L}[a, b]$ such that $v \in \text{Close}(t)$. For $v \in \{a, b\}$ the inequality is satisfied with equality; thus, we only need to care about the case $v \in \mathcal{L}(a, b) \cap \text{Close}(t)$.

Notice that $v \in \Gamma_s$, but on the other hand $t \notin \Gamma_s$. Therefore, we can take x as the last intersection of $P_{v,t}$ with $\partial\Gamma_s$. We cannot have $x = v$, otherwise the edge connecting x with the next vertex on the path would lie outside of the noose since $v \in \mathcal{L}(a, b)$. Also, since $v \in \text{Close}(t)$, we get that $P_{v,t} \cap V(\mathcal{L}) = \{v\}$; hence, $P_{v,t}$ cannot intersect \mathcal{L} at any vertex other than v . Therefore, $x \in P_{a,s} \cup P_{b,s}$. Without loss of generality, assume that $x \in P_{a,s}$.

As x lies both on the shortest path from a to s , and the shortest path from v to t , we have the following equalities:

$$\text{dist}(v, t) = \text{dist}(v, x) + \text{dist}(x, t) \quad \text{and} \quad \text{dist}(a, s) = \text{dist}(a, x) + \text{dist}(x, s).$$

Moreover, $\text{dist}(a, t) = \text{dist}(a, s) = d_a$.

From the triangle inequality and the equations above, we get

$$\text{dist}(a, x) + \text{dist}(x, t) \geq \text{dist}(a, t) = \text{dist}(a, s) = \text{dist}(a, x) + \text{dist}(x, s),$$

or equivalently, $\text{dist}(x, t) \geq \text{dist}(x, s)$.

Then, from the conditions above and another use of triangle inequality, we have

$$\text{dist}(v, t) = \text{dist}(v, x) + \text{dist}(x, t) \geq \text{dist}(v, x) + \text{dist}(x, s) \geq \text{dist}(v, s).$$

□

4.6 Conclusion of the proof of Noose Profile Lemma

Assume there are M vertices $s_1, s_2, \dots, s_M \in \mathcal{S}$ with pairwise distinct distance- d profiles on \mathcal{L} . Pick any vertices $s_i, s_j, i \neq j$, and let us describe their relation with respect to the regions $\Gamma_{s_i}, \Gamma_{s_j}$ defined by them:

- If $s_i \notin \Gamma_{s_j}$ and $s_j \notin \Gamma_{s_i}$, then from Lemma 48 we get $\pi_d[s_i, \mathcal{L}] = \pi_d[s_j, \mathcal{L}]$ — a contradiction.
- If $s_i \in \Gamma_{s_j}$ and $s_j \in \Gamma_{s_i}$, then from Lemma 45 we get $\Gamma_{s_i} \subsetneq \Gamma_{s_j} \subsetneq \Gamma_{s_i}$ — another contradiction.
- Hence, either $\pi_d[s_i, \mathcal{L}] \prec \pi_d[s_j, \mathcal{L}]$ (if $s_j \notin \Gamma_{s_i}$) or $\pi_d[s_j, \mathcal{L}] \prec \pi_d[s_i, \mathcal{L}]$ (if $s_i \notin \Gamma_{s_j}$).

Therefore, the relation \preccurlyeq defined on the distance- d profiles for s_1, s_2, \dots, s_M is antisymmetric, transitive (by definition of \preccurlyeq) and connex (by our considerations above). Hence, it is a linear order and thus contains a chain of length M . However, as π_d is a function whose domain contains c elements, and its codomain contains $d + 2$ possible values, the maximum length of such a chain is $c(d + 1) + 1$. We thus get $M \leq c(d + 1) + 1$.

Summing everything up, we proved that each set $\mathcal{S}(a, b, d_a, d_b)$ generates at most $c(d + 1) + 1$ different distance- d profiles. There are at most $c(c - 1)d^2$ such sets, as $a \neq b$ and $1 \leq d_a, d_b \leq d$. Hence, the number of different profiles induced in G by V_1 is bounded by

$$c(c - 1)d^2 \cdot [c(d + 1) + 1] \leq c^3(d + 1)^3.$$

The number of profiles in $\mathcal{C}_2 \cup \mathcal{C}_3 \cup \dots \cup \mathcal{C}_c = V_1 \cup V_2$ is then bounded by $c^3(d + 1)^4$ (Corollary 38). The vertices from $\mathcal{C}_0 \cup \mathcal{C}_1$ induce at most $c(d + 1) + 1$ additional profiles. The total number of profiles can be thus bounded by

$$c^3(d + 1)^4 + c(d + 1) + 1 \leq c^3(d + 2)^4.$$

This concludes the proof of the Noose Profile Lemma. We note a simple corollary which is a direct consequence of it:

Corollary 49. *Consider any graph G together with its embedding in the plane, and fix any noose \mathcal{L} containing $c \geq 1$ vertices of G . The set $\{\pi_d[v, \mathcal{L}] \mid v \notin V(G_{\mathcal{L}})\}$ of different distance- d profiles on \mathcal{L} , measured from the vertices **not enclosed** by \mathcal{L} , has at most $c^3(d + 2)^4$ elements.*

Proof. Let us take the embedding of G in the plane, and apply to it an inversion with any positive radius and a center in any point enclosed by \mathcal{L} that does not lie on any vertex or edge of G . After this transformation, the images of all vertices of G that have not been enclosed by \mathcal{L} in the original embedding will now be enclosed by the inversive

image \mathcal{L}' of the noose. Noose Profile Lemma applies to the inversed embedding as well, which allows us to reason that there are at most $c^3(d+2)^4$ different distance- d profiles on \mathcal{L} , measured from the vertices whose images are enclosed by \mathcal{L}' — or, in other words, measured from the vertices which were not enclosed by \mathcal{L} in the original embedding. \square

4.7 Proof of Theorem 26

We now lift the Noose Profile Lemma to the general setting, where the noose can be replaced with any non-empty set of vertices. For convenience, we restate the theorem here.

Theorem 26. *Consider any planar graph G , fix a non-empty set A containing $c \geq 1$ vertices of G . Then, the set $\{\pi_d[v, A] \mid v \in V(G)\}$ of all distance- d profiles on A has at most $64c^3(d+2)^7$ distinct elements.*

Proof. We fix G and embed it in the plane. We also fix the set A as in the statement of the theorem. Let $V_d = \{v \in V(G) \mid \exists_{s \in A} \text{dist}(v, s) \leq d\}$ be the set of vertices of G which are at distance at most d from any vertex in A . We observe that for each vertex $v \notin V_d$, the distance- d profile $\pi_d[v, A]$ is the constant function equal to $+\infty$. Now, we define the graph $G_d = G[V_d]$ induced by V_d . We can easily see that for each vertex $v \in V_d$, the distance- d profile of v on A in G_d is identical to the distance- d profile of v on A in G . Hence, the number of different distance- d profiles on A in G_d is at most 1 less than the number of different distance- d profiles on A in G .

We first assume that G_d is connected; we will resolve the case where G_d is disconnected at the end of the proof. Let \mathcal{T} be the minimum Steiner tree on A in G_d ; that is, the smallest possible tree which is a connected subgraph of G_d and spans all vertices of A .

Claim 50. \mathcal{T} has at most $(c-1)(2d+1)$ edges.

Proof. Let G_A be the complete weighted graph whose set of vertices is A , and for every pair of vertices $u, v \in A$, the edge between these vertices has weight $\text{dist}_{G_d}(u, v)$. Let also \mathcal{T}_A be the minimum-weight spanning tree of G_A .

We now prove that no edge of \mathcal{T}_A has weight greater than $2d+1$. Assume for contradiction that $uv \in E(\mathcal{T}_A)$ is an edge with weight at least $2d+2$. In this case, take a shortest path from u to v in G_d : $(x_0, x_1, x_2, \dots, x_\delta)$, where $x_0 = u$, $x_\delta = v$, $\delta \geq 2d+2$. We have that $\text{dist}_{G_d}(x_{d+1}, u) = d+1$ and $\text{dist}_{G_d}(x_{d+1}, v) = \delta - (d+1)$. As $x_{d+1} \in V_d$, there exists a vertex $s \in A$ at distance at most d from x_{d+1} . By the triangle inequality, we infer that

$$\begin{aligned} \text{dist}_{G_d}(s, u) &\leq \text{dist}_{G_d}(s, x_{d+1}) + \text{dist}_{G_d}(x_{d+1}, u) \leq d + (d+1) < \delta, \\ \text{dist}_{G_d}(s, v) &\leq \text{dist}_{G_d}(s, x_{d+1}) + \text{dist}_{G_d}(x_{d+1}, v) \leq d + (\delta - (d+1)) < \delta. \end{aligned}$$

Hence, G_A contains a three-vertex cycle (u, v, s) , in which the edge uv is strictly the heaviest. Therefore, uv cannot belong to \mathcal{T}_A — a contradiction.

Now, for each edge $uv \in \mathcal{T}_A$, we connect u with v in G_d using the shortest path between these two vertices. We can easily see that the union of these paths is a connected subgraph of G_d containing all the vertices in A . Hence, \mathcal{T} , the minimum Steiner tree on A , is at most as long as the union of these paths. Since we used $|A| - 1 = c - 1$ paths, and each of them has length at most $2d+1$, we infer that \mathcal{T} has at most $(c-1)(2d+1)$ edges. \square

We now create a new graph G_{cut} , which is a modified version of G_d , by “cutting the plane open” along the Steiner tree \mathcal{T} (Figure 17). This process is inspired by a similar idea by Pilipczuk et al. [32]. Formally, G_{cut} is the result of the following process:

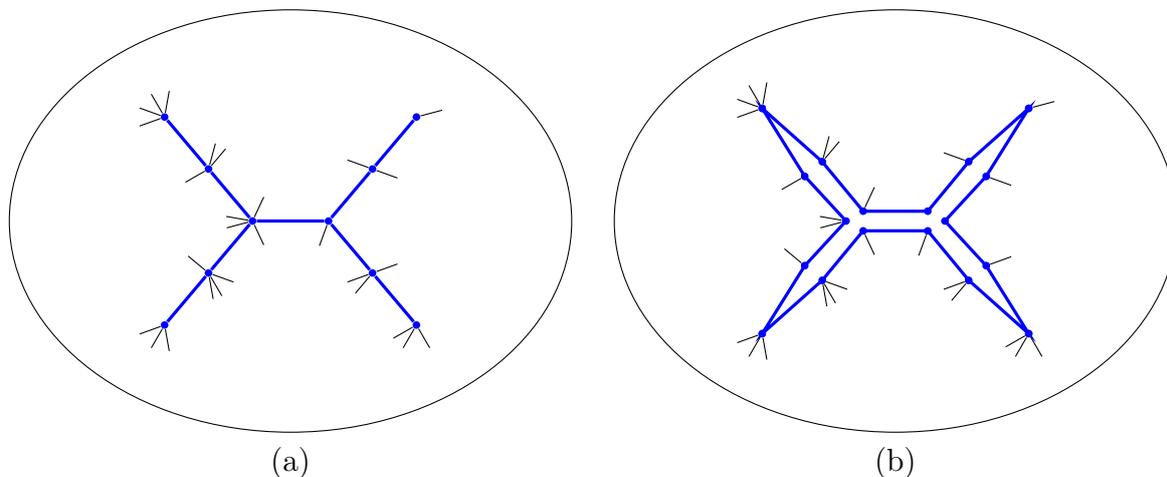


Figure 17: Creating G_{cut} . (a) — original graph G_d , with tree \mathcal{T} (blue). (b) — the resulting graph, with an Euler tour of \mathcal{T} .

1. We find an Euler tour of \mathcal{T} which does not intersect itself in the plane embedding of G_d . This tour traverses each edge of \mathcal{T} twice, in the opposite directions.
2. For each vertex $v \in V(\mathcal{T})$, we let $s_1, s_2, \dots, s_\Delta$ to be the neighbors of v in \mathcal{T} , ordered anti-clockwise around v . For simplicity, we assume $s_{\Delta+1} = s_1$. If $\Delta \geq 2$, we create Δ copies of v (including the original vertex v): $v_1, v_2, \dots, v_\Delta$, and we put the i -th copy, v_i , very close to the original vertex v , and between the edges vs_i and vs_{i+1} in the embedding of G_d . Moreover, for each neighbor x of v such that $vx \notin E(\mathcal{T})$, we find the only index i such that the edges vs_i, vx, vs_{i+1} are ordered anti-clockwise around v ; then, we replace the edge vx with $v_i x$.
3. We duplicate each edge uv of the Euler tour of \mathcal{T} , so that we have one copy for each orientation of uv in the Euler tour. We then use both copies to connect the corresponding copies of vertices u, v in G_{cut} . We can easily see that all new edges form a cycle enclosing \mathcal{T} , which we call C . Naturally, C has twice as many edges as \mathcal{T} .
4. We remove the edges of \mathcal{T} from the graph.

Let $V(C)$ be the set of vertices of C , and $E(C)$ — the set of edges of C . Since $|E(\mathcal{T})| \leq (c-1)(2d+1)$ (Claim 50), we have that $|E(C)| = 2|E(\mathcal{T})| \leq 2(c-1)(2d+1)$. Also, for every vertex $v \in V(\mathcal{T})$, we let $\deg_{\mathcal{T}}(v)$ to be the degree of v in \mathcal{T} . Hence, for $v \in V(\mathcal{T})$, the graph G_{cut} contains vertices $v_1, \dots, v_{\deg_{\mathcal{T}}(v)}$.

Claim 51. *For any two different vertices $p, q \in V(G_{\text{cut}}) - V(C)$, if $\pi_d[p, V(C)] = \pi_d[q, V(C)]$ in G_{cut} , then $\pi_d[p, V(\mathcal{T})] = \pi_d[q, V(\mathcal{T})]$ in G_d .*

Proof. We will show that the distance- d profile of any vertex $p \in V(G_{\text{cut}}) - V(C)$ on $V(C)$ in G_{cut} uniquely determines the distance- d profile of p on $V(\mathcal{T})$ in G_d . The statement of the lemma will follow.

We first remark that for each $x \in V(\mathcal{T})$ and $i \in [1, \deg_{\mathcal{T}}(x)]$, we have $\text{dist}_{G_{\text{cut}}}(p, x_i) \geq \text{dist}_{G_d}(p, x)$. This is true since we can pick a shortest path from p to x_i in G_{cut} , and replace each vertex of this path belonging to $V(C)$ with the corresponding vertex of $V(\mathcal{T})$ in G_d . The resulting path connects p with x in G_d , and has length $\text{dist}_{G_{\text{cut}}}(p, x_i)$.

We fix a vertex $v \in V(\mathcal{T})$, and we want to determine $\pi_d[p, V(\mathcal{T})](v)$ in G_d from $\pi_d[p, V(C)]$ in G_{cut} . Obviously, for each $x \in V(\mathcal{T})$ and $i \in [1, \deg_{\mathcal{T}}(x)]$ we have

$$\text{dist}_{G_d}(p, v) \leq \text{dist}_{G_d}(p, x) + \text{dist}_{G_d}(x, v) \leq \text{dist}_{G_{\text{cut}}}(p, x_i) + \text{dist}_{G_d}(x, v).$$

On the other hand, we pick a shortest path P from p to v in G_d . We let x to be the first intersection of P with $V(\mathcal{T})$, and y — the vertex immediately preceding x on P . Since $y \notin V(\mathcal{T})$, then by the construction of G_{cut} we infer that y is adjacent to a copy of x in G_{cut} — say, x_i for some $i \in [1, \deg_{\mathcal{T}}(x)]$. For this choice of x and i we have that

$$\begin{aligned} \text{dist}_{G_d}(p, v) &= \text{dist}_{G_d}(p, y) + 1 + \text{dist}_{G_d}(x, v) = \\ &= \text{dist}_{G_{\text{cut}}}(p, y) + \text{dist}_{G_{\text{cut}}}(y, x_i) + \text{dist}_{G_d}(x, v) \geq \\ &\geq \text{dist}_{G_{\text{cut}}}(p, x_i) + \text{dist}_{G_d}(x, v). \end{aligned}$$

We conclude that

$$\text{dist}_{G_d}(p, v) = \min\{\text{dist}_{G_{\text{cut}}}(p, x_i) + \text{dist}_{G_d}(x, v) \mid x \in V(\mathcal{T}), i \in [1, \deg_{\mathcal{T}}(x)]\}.$$

It is now straightforward to determine $\pi_d[p, V(\mathcal{T})](v)$ in G_d from $\pi_d[p, V(C)]$ in G_{cut} and the distances in G_d between each pair of vertices of \mathcal{T} (which are independent on the choice of p). \square

We now observe that C is a cycle in G_{cut} , which does not enclose any other vertices of G_{cut} ; hence, there exists a noose \mathcal{L} , closely following C , that passes through each vertex of C exactly once and that does not enclose any other vertex of G_{cut} . The variant of Noose Profile Lemma (Corollary 49) applies to \mathcal{L} . Hence, there are at most $|C|^3(d+2)^4$ different distance- d profiles on C in G_{cut} , measured from the vertices not enclosed by \mathcal{L} .

We fix two vertices $p, q \in V(G_{\text{cut}}) - V(C)$ such that $\pi_d[p, V(C)] = \pi_d[q, V(C)]$ in G_{cut} . Firstly, we see that neither p nor q are enclosed by \mathcal{L} . We know from Claim 51 that $\pi_d[p, V(\mathcal{T})] = \pi_d[q, V(\mathcal{T})]$ in G_d . Since $A \subseteq V(\mathcal{T})$, we infer that $\pi_d[p, A] = \pi_d[q, A]$ in G_d . By the construction of G_d , we also get that $\pi_d[p, A] = \pi_d[q, A]$ in G .

Since there are at most $|C|^3(d+2)^4$ different distance- d profiles on C in G_{cut} , measured from the vertices not enclosed by \mathcal{L} , we get that there are at most $|C|^3(d+2)^4$ different distance- d profiles on A in G_d , measured from the vertices in $V(G_d) - V(\mathcal{T})$. Therefore, G_d has at most $|C|^3(d+2)^4 + |V(\mathcal{T})|$ different distance- d profiles on A in total. Recall that G has at most one more distance- d profile on A compared to G_d — that is, the constant function equal to $+\infty$. Hence, the number of distance- d profiles on A in G is bounded by

$$\begin{aligned} |C|^3(d+2)^4 + |V(\mathcal{T})| + 1 &\leq [2(c-1)(2d+1)]^3(d+2)^4 + (c-1)(2d+1) + 1 \leq \\ &\leq 8c^3[(2d+1)^3(d+2)^4 + (2d+1) + 1] \leq 64c^3(d+2)^7. \end{aligned}$$

Let us finally consider the case where G_d is disconnected. Assume that the vertices of A are spread among k different connected components in G_d , and the i -th connected component contains c_i vertices of A for each $i \in [1, k]$ ($c_1, c_2, \dots, c_k \geq 1$, $\sum c_i = c$). Then, we apply this theorem to each connected component of G_d separately. We get that the number of different distance- d profiles on A in G_d is bounded by

$$\sum_{i=1}^k 64c_i^3(d+2)^7 < 64c^3(d+2)^7.$$

Hence, the number of different distance- d profiles on A in G is also bounded by $64c^3(d+2)^7$. \square

5 Upper bound on the orders of semi-ladders in planar graphs

This section is fully dedicated to the proof of an upper bound on the maximum order of a distance- d semi-ladder in planar graphs:

Theorem 52. *The order of any distance- d semi-ladder in a planar graph is bounded from above by $d^{O(d)}$.*

We refer to the introduction (Section 1) for a rough sketch of the proof. Here, we provide more details on the structure of the proof.

In Subsection 5.1, we introduce two structural objects: quasi-cages and cages. These configurations occur in planar graphs with sufficiently large semi-ladders. Namely, we prove in Lemma 55 that each graph containing a distance- d semi-ladder of order $d(dl + 2)^d + 1$ also contains a quasi-cage of order ℓ . Then, we establish in Lemma 61 that each graph containing a quasi-cage of order $\Theta(d) \cdot \ell$ also contains a cage of order ℓ .

In Subsection 5.2, we prove multiple structural properties of cages. These properties allow us to introduce two variants of a cage: ordered cage and identity ordered cage, which are structures utilizing the topology of the embedding of the graph in the plane. We then prove that each graph with a cage of order ℓ also contains an ordered cage of order ℓ (Lemma 67), and that each graph with an ordered cage of order $\Theta(\ell^2)$ contains an identity ordered cage of order ℓ (Lemma 73).

In Subsections 5.3 and 5.4, we find more structural properties of identity ordered cages, building on prior knowledge about cages. Thanks to these properties, we establish another variant of a cage — neighbor cage — and we prove that each graph with an identity ordered cage of order $\Theta(d^7) \cdot \ell$ also contains a neighbor cage of order ℓ (Lemma 84). The polynomial bound on neighborhood complexity in planar graphs becomes essential in this part of the proof. We then proceed to prove a set of properties of neighbor cages.

Finally, in Subsection 5.5, we introduce the final variant of a cage — separating cage — and prove that if a graph contains a neighbor cage of order $\Theta(d^2) \cdot \ell$, then it also contains a separating cage of order ℓ (Lemma 95). We finalize the proof by directly showing that each separating cage has order smaller than $2d + 5$ (Lemma 97). Retracing all the steps, we can easily figure out the $d^{O(d)}$ upper bound on the maximum order of a distance- d semi-ladder in the class of planar graphs (Corollary 98).

Throughout this section, we consider a fixed planar graph G together with its embedding in the plane; the embedding may be altered later in the proof. We also fix the parameter $d \geq 1$.

5.1 Quasi-cages and cages

We begin by introducing a prototype for cages — quasi-cages. In order to do this, we first need to define a special kind of a rooted tree.

Definition 53. Paths P_1, P_2, \dots, P_ℓ in a planar graph G create a **geodesic tree** $\mathcal{T} = \bigcup P_i = \text{Tree}(P_1, P_2, \dots, P_\ell)$ of order ℓ rooted at a vertex r if the following conditions hold:

- the union of all the paths is a subtree \mathcal{T} of G ,
- all paths P_i are shortest paths in G and have the same length, not exceeding d ,
- each path P_i originates at r and terminates at a different leaf of the tree.

We say that \mathcal{T} **avoids** s if none of the paths contain s as a vertex. Also, if for all different i, j , the intersection of P_i and P_j is equal to $\{r\}$, we call \mathcal{T} a **simple geodesic tree**.

We also use the following notation: $\text{Root}(\mathcal{T})$ is the root r of the geodesic tree, and $\text{Path}(\mathcal{T}, i)$ is the path P_i . Moreover, for any two vertices $u, v \in \mathcal{T}$, we define $\text{lca}_{\mathcal{T}}(u, v)$ as the lowest common ancestor of these two vertices in \mathcal{T} , where \mathcal{T} is considered rooted at r . We may write lca instead of $\text{lca}_{\mathcal{T}}$ if the tree is clear from the context.

We can now define a quasi-cage:

Definition 54. A **quasi-cage** of order ℓ in a planar graph G is a structure consisting of:

- a semi-ladder of order ℓ consisting of distinct vertices a_1, a_2, \dots, a_ℓ , and b_1, b_2, \dots, b_ℓ ,
- two distinct root vertices $p, q \notin \{a_1, a_2, \dots, a_\ell, b_1, b_2, \dots, b_\ell\}$,
- paths P_1, P_2, \dots, P_ℓ , where P_i is a shortest path from p to a_i for each $i \in [1, \ell]$, which form a simple geodesic tree $\mathcal{P} = \text{Tree}(P_1, P_2, \dots, P_\ell)$ rooted at p and avoiding q ,
- paths Q_1, Q_2, \dots, Q_ℓ , where Q_i is a shortest path from q to a_i for each $i \in [1, \ell]$, which form a geodesic tree $\mathcal{Q} = \text{Tree}(Q_1, Q_2, \dots, Q_\ell)$ rooted at q and avoiding p .

We stress that while \mathcal{P} must be a simple geodesic tree, \mathcal{Q} is not necessarily simple (Figure 18).

With the notation above, we denote a quasi-cage by

$$\text{QuasiCage}((a_1, a_2, \dots, a_\ell), (b_1, b_2, \dots, b_\ell), \mathcal{P}, \mathcal{Q}).$$

We remark that the definitions of semi-ladders, geodesic trees, simple geodesic trees, and quasi-cages are closed to taking subsets. Formally, fix a subsequence of indices i_1, i_2, \dots, i_m ($1 \leq i_1 < i_2 < \dots < i_m \leq \ell$). Then:

- For a semi-ladder $(a_1, a_2, \dots, a_\ell), (b_1, b_2, \dots, b_\ell)$, we can easily see that vertices $(a_{i_1}, a_{i_2}, \dots, a_{i_m}), (b_{i_1}, b_{i_2}, \dots, b_{i_m})$ form a semi-ladder of order m .

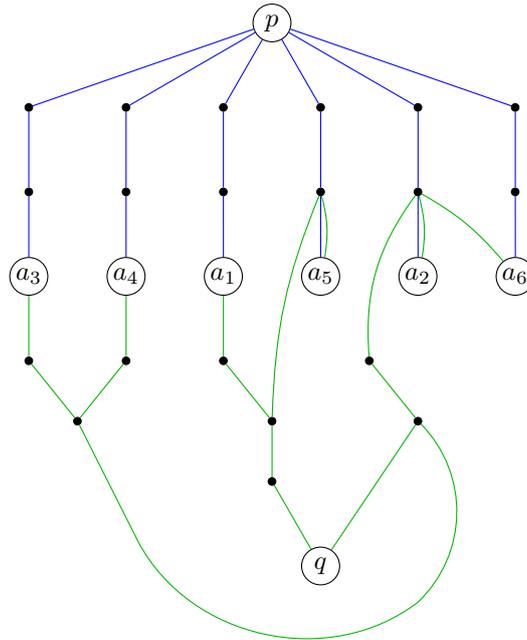


Figure 18: An example quasi-cage of order 6. Tree \mathcal{P} is marked blue and tree \mathcal{Q} is marked green. Vertices b_1, b_2, \dots, b_6 were omitted from the figure. Note that trees \mathcal{P} and \mathcal{Q} can share vertices and edges, but \mathcal{P} must not contain q , and \mathcal{Q} must not contain p .

- For a geodesic tree $\text{Tree}(P_1, P_2, \dots, P_\ell)$ of order ℓ , also $\text{Tree}(P_{i_1}, P_{i_2}, \dots, P_{i_m})$ is a geodesic tree of order m . If the original geodesic tree was simple, its subset is simple as well.

- For a quasi-cage

$$\text{QuasiCage}((a_1, a_2, \dots, a_\ell), (b_1, b_2, \dots, b_\ell), \text{Tree}(P_1, P_2, \dots, P_\ell), \text{Tree}(Q_1, Q_2, \dots, Q_\ell))$$

of order ℓ , the following describes a quasi-cage of order m :

$$\text{QuasiCage}((a_{i_1}, \dots, a_{i_m}), (b_{i_1}, \dots, b_{i_m}), \text{Tree}(P_{i_1}, \dots, P_{i_m}), \text{Tree}(Q_{i_1}, \dots, Q_{i_m})).$$

This fact will also be true for each further refinement of the definition of a quasi-cage. Therefore, from now on, we will refer to the subsets of any configuration (semi-ladder, geodesic tree, quasi-cage and any its refinement) by taking the original configuration and specifying the indices i_1, i_2, \dots, i_m .

As the first step of our proof, we will demonstrate that planar graphs with large distance- d semi-ladders also contain large quasi-cages.

Lemma 55. *If G contains a distance- d semi-ladder of order $d(dl + 2)^d + 1$, then G also contains a quasi-cage of order ℓ .*

Proof. We will state and prove a series of claims, each imposing more structure on the semi-ladder.

We will call a distance- d semi-ladder $(a_1, a_2, \dots, a_k), (b_1, b_2, \dots, b_k)$ of order k , together with another vertex r , **R-equidistant** if:

- $r \notin \{a_1, a_2, \dots, a_k, b_1, b_2, \dots, b_k\}$,
- each vertex a_i is at the same distance from r , not exceeding d .

We will denote such semi-ladder as $\text{RSemiLadder}((a_1, \dots, a_k), (b_1, \dots, b_k), r)$.

Claim 56. *Every distance- d semi-ladder of order $dl + 1$ contains an R -equidistant semi-ladder of order ℓ as a subset.*

Proof. Fix a distance- d semi-ladder of order $dl + 1$: $(a_1, a_2, \dots, a_{dl+1}), (b_1, b_2, \dots, b_{dl+1})$. Notice that $1 \leq \text{dist}(a_i, b_1) \leq d$ for each $i \geq 2$. This means that the set $\{a_2, a_3, \dots, a_{dl+1}\}$ of vertices contains a subset with at least $\frac{(dl+1)-1}{d} = \ell$ elements where each element is at the same distance from b_1 , not exceeding d . Let this subset be $a_{i_1}, a_{i_2}, \dots, a_{i_\ell}$ ($2 \leq i_1 < i_2 < \dots < i_\ell \leq dl + 1$). Then, the following structure is an R -equidistant semi-ladder:

$$\text{RSemiLadder}((a_{i_1}, \dots, a_{i_\ell}), (b_{i_1}, \dots, b_{i_\ell}), b_1).$$

By our definition of a semi-ladder, the vertices of any semi-ladder must be different. Hence, $b_1 \notin \{a_{i_1}, \dots, a_{i_\ell}, b_{i_1}, \dots, b_{i_\ell}\}$. \lrcorner

We define a **simple geodesic semi-ladder** of order k as a distance- d semi-ladder $(a_1, a_2, \dots, a_k), (b_1, b_2, \dots, b_k)$, together with a simple geodesic tree $\mathcal{P} = \text{Tree}(P_1, \dots, P_k)$, where for each $i \in [1, k]$, the path P_i terminates at a_i . We denote such a simple geodesic semi-ladder as $\text{SGSemiLadder}((a_1, \dots, a_k), (b_1, \dots, b_k), \mathcal{P})$.

Claim 57. *Every R -equidistant semi-ladder of order ℓ^d contains a simple geodesic semi-ladder of order ℓ as a subset.*

Proof. Take an R -equidistant semi-ladder of order ℓ^d :

$$\text{RSemiLadder}((a_1, \dots, a_{\ell^d}), (b_1, \dots, b_{\ell^d}), r).$$

Let $\delta := \text{dist}(r, a_1)$. We also take \mathcal{T} as a shortest paths tree rooted at r whose leaves are exactly $a_1, a_2, \dots, a_{\ell^d}$; this tree exists since all hypothetical leaves are at the same distance $\delta \leq d$ from r . This tree has depth δ and ℓ^d leaves. Thus, \mathcal{T} contains a vertex which has at least $\ell^{d/\delta} \geq \ell$ children. Let p be any such vertex.

Since p has at least ℓ children in \mathcal{T} , and each child is a root of a subtree in \mathcal{T} containing at least one leaf of \mathcal{T} , we can pick a subset $A = \{a_{i_1}, a_{i_2}, \dots, a_{i_\ell}\}$ ($1 \leq i_1 < i_2 < \dots < i_\ell \leq \ell^d$) of leaves of \mathcal{T} where each leaf is in a different subtree of \mathcal{T} rooted at a child of p . Also, for each $j \in [1, \ell]$, we define P_j as the shortest path in \mathcal{T} which originates at p and terminates at a_{i_j} . By the choice of A , all paths P_1, P_2, \dots, P_ℓ are vertex-disjoint apart from their common origin p . They also have the same length since all leaves of \mathcal{T} are at the same depth, and all leaves in A are in the subtree of \mathcal{T} rooted at p . Finally, they are shortest paths in G by the choice of \mathcal{T} as a shortest paths tree. Therefore, $\text{Tree}(P_1, P_2, \dots, P_\ell)$ is a simple geodesic tree.

We can now see that the following structure is a simple geodesic semi-ladder:

$$\text{SGSemiLadder}((a_{i_1}, \dots, a_{i_\ell}), (b_{i_1}, \dots, b_{i_\ell}), \text{Tree}(P_1, P_2, \dots, P_\ell)).$$

\lrcorner

A simple geodesic semi-ladder $\text{SGSemiLadder}((a_1, \dots, a_k), (b_1, \dots, b_k), \mathcal{P})$ of order k , together with another vertex q , is **Q-equidistant** if:

- \mathcal{P} avoids q ,
- $q \notin \{a_1, a_2, \dots, a_k, b_1, b_2, \dots, b_k\}$,
- q is at the same distance from each vertex a_i , and
- $\text{dist}(q, \text{Root}(\mathcal{P})) + \text{dist}(\text{Root}(\mathcal{P}), a_i) > d$ for each $i \in [1, k]$.

We use the following notation for a Q-equidistant simple geodesic semi-ladder:

$$\text{QSemiLadder}((a_1, \dots, a_k), (b_1, \dots, b_k), \mathcal{P}, q).$$

Claim 58. *Every simple geodesic semi-ladder of order $d\ell + 2$ contains a Q-equidistant simple geodesic semi-ladder of order ℓ as a subset.*

Proof. Take a simple geodesic semi-ladder of order $d\ell + 2$:

$$\text{SGSemiLadder}((a_1, \dots, a_{d\ell+2}), (b_1, \dots, b_{d\ell+2}), \mathcal{P}) \quad \text{where } \mathcal{P} = \text{Tree}(P_1, P_2, \dots, P_{d\ell+2}).$$

Notice that $b_1 \neq \text{Root}(\mathcal{P})$; otherwise, we would have

$$d < \text{dist}(b_1, a_1) = \text{dist}(\text{Root}(\mathcal{P}), a_1) = \text{dist}(\text{Root}(\mathcal{P}), a_2) = \text{dist}(b_1, a_2) \leq d.$$

Since paths $P_2, \dots, P_{d\ell+2}$ are vertex-disjoint apart from $\text{Root}(\mathcal{P})$, this means that at most one of them passes through b_1 . Hence, at least $d\ell$ such paths do not contain b_1 .

Moreover, for each $i \in [2, d\ell + 2]$, we have $\text{dist}(a_i, b_1) \in [1, d]$. By the pigeonhole principle, we can fix ℓ indices i_1, i_2, \dots, i_ℓ ($2 \leq i_1 < i_2 < \dots < i_\ell \leq d\ell + 2$) with the following properties:

- for each $j \in [1, \ell]$, the path P_{i_j} does not contain b_1 ,
- b_1 is at the same distance from each vertex a_{i_j} ($j \in [1, \ell]$).

We will prove that the following structure is a Q-equidistant simple geodesic semi-ladder:

$$\text{QSemiLadder}((a_{i_1}, \dots, a_{i_\ell}), (b_{i_1}, \dots, b_{i_\ell}), \text{Tree}(P_{i_1}, \dots, P_{i_\ell}), b_1).$$

Since the choice of the indices i_1, \dots, i_ℓ already met the first and the third condition in the definition of a Q-equidistant simple geodesic semi-ladder, we only need to verify the remaining requirements.

As any two vertices in the semi-ladder are different, we have $b_1 \neq a_{i_j}$ and $b_1 \neq b_{i_j}$ for each $j \in [1, \ell]$, which confirms the second condition. Moreover, for every $j \in [1, \ell]$, we have that

$$\begin{aligned} \text{dist}(b_1, \text{Root}(\mathcal{P})) + \text{dist}(\text{Root}(\mathcal{P}), a_{i_j}) &= \text{dist}(b_1, \text{Root}(\mathcal{P})) + \text{dist}(\text{Root}(\mathcal{P}), a_1) \geq \\ &\geq \text{dist}(b_1, a_1) > d. \end{aligned}$$

All the required conditions are satisfied, so the structure defined above is a Q-equidistant simple geodesic semi-ladder. ┘

Now we will see that a large Q-equidistant simple geodesic semi-ladder in G directly witnesses the existence of a large quasi-cage in G .

Claim 59. *Every graph with a Q-equidistant simple geodesic semi-ladder of order ℓ also contains a quasi-cage of order ℓ .*

Proof. Fix a Q-equidistant simple geodesic semi-ladder as in the statement of the claim:

$$\mathcal{L} = \text{QSemiLadder}((a_1, \dots, a_\ell), (b_1, \dots, b_\ell), \mathcal{P}, q).$$

Let also \mathcal{S} denote the semi-ladder $a_1, \dots, a_\ell, b_1, \dots, b_\ell$ underlying \mathcal{L} , and set $p := \text{Root}(\mathcal{P})$. Let us also take \mathcal{Q} as any shortest-path tree rooted at q whose leaves are a_1, a_2, \dots, a_ℓ .

We will now verify that according to Definition 54, the following structure is a quasi-cage:

$$\text{QuasiCage}((a_1, \dots, a_\ell), (b_1, \dots, b_\ell), \mathcal{P}, \mathcal{Q}).$$

We first note that $p \notin \mathcal{S}$ since \mathcal{L} is also an R-equidistant semi-ladder. Similarly, q is disjoint with \mathcal{S} , which follows immediately from the definition of \mathcal{L} as a Q-equidistant simple geodesic semi-ladder. The geodesic tree \mathcal{P} is simple (since \mathcal{L} is a simple geodesic semi-ladder) and avoids q (since \mathcal{L} is a Q-equidistant simple geodesic semi-ladder).

We only need to prove that \mathcal{Q} avoids p . Assume for contradiction that for some $i \in [1, \ell]$, a path $Q_i = \text{Path}(\mathcal{Q}, i)$ contains p as a vertex. This would mean that the shortest path between q and a_i contains p as a vertex. However, by the definition of a Q-equidistant simple geodesic semi-ladder:

$$\text{dist}(q, a_i) = \text{dist}(q, p) + \text{dist}(p, a_i) = \text{dist}(q, \text{Root}(\mathcal{P})) + \text{dist}(\text{Root}(\mathcal{P}), a_i) > d$$

— a contradiction. ┘

We can now combine the claims above to prove our lemma. Take any distance- d semi-ladder L_1 in G of order $d(d\ell + 2)^d + 1$.

By Claim 56, L_1 contains an R-equidistant semi-ladder L_2 of order $(d\ell + 2)^d$ as a subset.

By Claim 57, L_2 contains a simple geodesic semi-ladder L_3 of order $d\ell + 2$ as a subset.

By Claim 58, L_3 contains a Q-equidistant simple geodesic semi-ladder L_4 of order ℓ as a subset.

Hence, by Claim 59, L_4 asserts the existence of a quasi-cage of order ℓ in G . □

Next, we refine the definition of a quasi-cage by defining cages — variants of quasi-cages ensuring that both geodesic trees in a definition of a quasi-cage intersect in a controlled way.

Definition 60. A **cage** of order ℓ is a quasi-cage $\text{QuasiCage}((a_1, \dots, a_\ell), (b_1, \dots, b_\ell), \mathcal{P}, \mathcal{Q})$ in which for every two indices $i, j \in [1, \ell]$, the paths $\text{Path}(\mathcal{P}, i)$ and $\text{Path}(\mathcal{Q}, j)$ have a non-empty intersection if and only if $i = j$.

We denote this structure in the following way:

$$\text{Cage}((a_1, \dots, a_\ell), (b_1, \dots, b_\ell), \mathcal{P}, \mathcal{Q}).$$

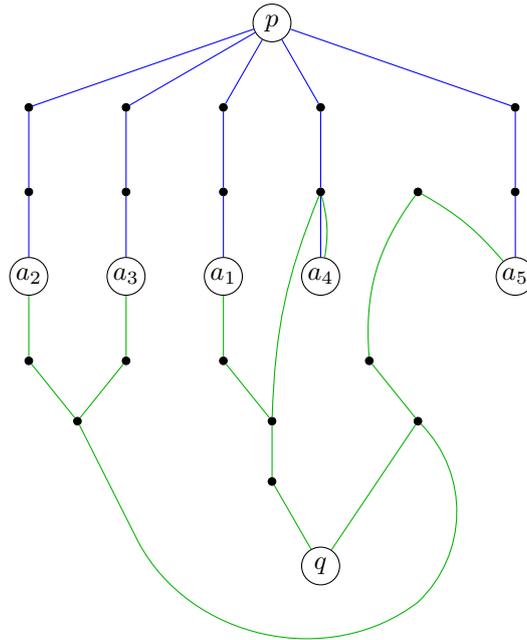


Figure 19: An example cage. This cage is a subset of the quasi-cage from Figure 18. The pair of vertices (a_2, b_2) has been removed from the original semi-ladder, and the remaining indices have been renumbered. Analogously to Figure 18, vertices b_1, b_2, \dots, b_5 were removed from the picture.

To see that the definition of a cage is slightly stronger than the one of a quasi-cage, we can verify that a quasi-cage in Figure 18 is not a cage — its path connecting q with a_6 intersects the path connecting p with a_2 . However, a large subset of it is a cage (Figure 19). This is not a coincidence, as proved in the following lemma.

Lemma 61. *Every graph with a quasi-cage of order $(2d - 1)\ell$ contains a cage of order ℓ as a subset.*

Proof. Fix a quasi-cage:

$$\text{QuasiCage}((a_1, a_2, \dots, a_{(2d-1)\ell}), (b_1, b_2, \dots, b_{(2d-1)\ell}), \mathcal{P}, \mathcal{Q})$$

where $\mathcal{P} = \text{Tree}(P_1, \dots, P_{(2d-1)\ell})$, $\mathcal{Q} = \text{Tree}(Q_1, \dots, Q_{(2d-1)\ell})$.

Firstly, note that for every $i \in [1, (2d - 1)\ell]$, paths P_i and Q_i intersect, because they share their final vertex a_i . Our aim is now to find a subset of indices $I \subseteq \{1, 2, \dots, (2d - 1)\ell\}$ such that $|I| = \ell$ and for any two different indices $i, j \in I$, paths P_i and Q_j do not intersect. As soon as we achieve this, the subset of the quasi-cage given by the set of indices I will form the sought cage.

Let us create an auxiliary directed graph H where $V(H) = \{1, 2, \dots, (2d - 1)\ell\}$, such that H contains an oriented edge $i \rightarrow j$ for $i \neq j$ if and only if P_i and Q_j intersect. Let also \widehat{H} be the undirected graph underlying H . We plan to find an independent set $I \subseteq V(\widehat{H})$ of size at least ℓ . Then for any two indices $i, j \in I$, we will have that paths P_i and Q_j are disjoint, as well as paths P_j and Q_i . Therefore, I will be a subset of indices forming a cage.

We note that each path Q_j has length at most d . As Q_j does not pass through p (the root of the simple geodesic tree \mathcal{P}), each vertex other than its origin q can lie on at most one of the paths P_i . Thus, Q_j intersects at most d different paths in \mathcal{P} . Since Q_j intersects with P_j for every $j \in [1, (2d-1)\ell]$, we get that Q_j has a non-empty intersection with at most $d-1$ paths P_i for $i \neq j$. This means that no vertex of H has in-degree larger than $d-1$.

In every non-empty subset S of vertices of \widehat{H} , the total number of edges in the subgraph of \widehat{H} induced by S does not exceed $|S| \cdot (d-1)$; this is because in the subgraph of H induced by S , each vertex has in-degree not exceeding $d-1$. Thus, \mathcal{S} contains a vertex with degree not exceeding the average degree of the vertices in S , which can be bounded from above by

$$\frac{2 \cdot [|S| \cdot (d-1)]}{|S|} = 2d - 2.$$

Hence, \widehat{H} is $(2d-2)$ -degenerate [18]. Since $(2d-2)$ -degenerate graphs are $(2d-1)$ -colorable [41], \widehat{H} contains an independent set of size at least

$$\frac{|V(\widehat{H})|}{2d-1} = \ell. \quad \square$$

5.2 Ordered cages

So far, the definition of a cage does not involve any topological properties of the plane. In this subsection, we will introduce the topology to the cages by defining ordered cages. We will, however, need some notation beforehand.

Definition 62. For a cage $\mathcal{C} := \text{Cage}((a_1, \dots, a_\ell), (b_1, \dots, b_\ell), \mathcal{P}, \mathcal{Q})$ of order ℓ in a graph G , and an index $i \in \{1, 2, \dots, \ell\}$, we define:

- the **i -th splitting vertex** $\text{Split}(\mathcal{C}, i)$ as the vertex of the intersection of $\text{Path}(\mathcal{P}, i)$ and $\text{Path}(\mathcal{Q}, i)$ which is the closest to $\text{Root}(\mathcal{P})$ (this vertex is unique since $\text{Path}(\mathcal{P}, i)$ is a shortest path in G);
- the **i -th splitting path** $\text{SPath}(\mathcal{C}, i)$ as the simple path originating at $\text{Root}(\mathcal{P})$, following the prefix of $\text{Path}(\mathcal{P}, i)$ until $\text{Split}(\mathcal{C}, i)$, and then — the prefix of $\text{Path}(\mathcal{Q}, i)$ in the opposite direction. Formally,

$$\text{SPath}(\mathcal{C}, i) = \text{Path}(\mathcal{P}, i)[\text{Root}(\mathcal{P}), \text{Split}(\mathcal{C}, i)] \cdot (\text{Path}(\mathcal{Q}, i))^{-1}[\text{Split}(\mathcal{C}, i), \text{Root}(\mathcal{Q})].$$

The definitions are also explained in Figure 20.

We derive the following facts about splitting paths:

Lemma 63. Fix a cage $\mathcal{C} = ((a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, \mathcal{P}, \mathcal{Q})$, two different indices $i, j \in [1, \ell]$ and a vertex $x \in V(\mathcal{Q})$. If $x \in Q_i$ and $x \in Q_j$, then $x \in \text{SPath}(\mathcal{C}, i) \cap \text{SPath}(\mathcal{C}, j)$.

Proof. Note that $\text{Split}(\mathcal{C}, i) \in P_i \cap Q_i$. Hence, path Q_j cannot contain $\text{Split}(\mathcal{C}, i)$ or otherwise paths P_i and Q_j would intersect. Therefore, x is an ancestor of $\text{Split}(\mathcal{C}, i)$ in \mathcal{Q} , so

$$x \in (\text{Path}(\mathcal{Q}, i))^{-1}[\text{Split}(\mathcal{C}, i), \text{Root}(\mathcal{Q})].$$

Thus, $x \in \text{SPath}(\mathcal{C}, i)$.

We prove that $x \in \text{SPath}(\mathcal{C}, j)$ analogously. □

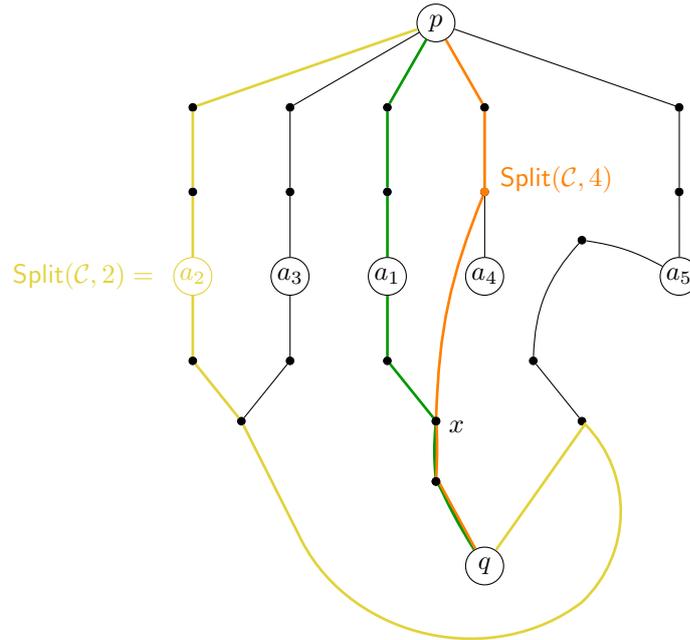


Figure 20: Example splitting vertices and splitting paths in the cage \mathcal{C} from Figure 19. Paths $\text{Split}(\mathcal{C}, 2)$ and $\text{SPath}(\mathcal{C}, 2)$ are marked yellow; $\text{Split}(\mathcal{C}, 4)$ and $\text{SPath}(\mathcal{C}, 4)$ are marked orange, and $\text{SPath}(\mathcal{C}, 1)$ is marked green.

Lemma 64. For a cage $\mathcal{C} = ((a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, \mathcal{P}, \mathcal{Q})$ in a graph G , and two different indices $i, j \in [1, \ell]$:

$$\text{SPath}(\mathcal{C}, i) \cap \text{SPath}(\mathcal{C}, j) = \{\text{Root}(\mathcal{P})\} \cup \text{SPath}(\mathcal{C}, i)[\text{lca}_{\mathcal{Q}}(a_i, a_j), \text{Root}(\mathcal{Q})].$$

In other words, any two distinct paths $\text{SPath}(\mathcal{C}, i)$, $\text{SPath}(\mathcal{C}, j)$ intersect exactly in their first vertex p and in their common suffix.

Proof. We first note that the following condition is simple to verify. By Lemma 63, we have that

$$\text{SPath}(\mathcal{C}, i)[\text{lca}_{\mathcal{Q}}(a_i, a_j), \text{Root}(\mathcal{Q})] = \text{SPath}(\mathcal{C}, j)[\text{lca}_{\mathcal{Q}}(a_i, a_j), \text{Root}(\mathcal{Q})].$$

Hence:

$$\text{SPath}(\mathcal{C}, i) \cap \text{SPath}(\mathcal{C}, j) \supseteq \{\text{Root}(\mathcal{P})\} \cup \text{SPath}(\mathcal{C}, i)[\text{lca}_{\mathcal{Q}}(a_i, a_j), \text{Root}(\mathcal{Q})].$$

We only need to prove the \subseteq part of the equality.

Let $\mathcal{P} = \text{Tree}(P_1, P_2, \dots, P_\ell)$ and $\mathcal{Q} = \text{Tree}(Q_1, Q_2, \dots, Q_\ell)$, and fix two different indices i, j . By the definition of a cage, paths P_i and Q_j are vertex-disjoint, and paths P_j and Q_i are vertex-disjoint. Since $\text{SPath}(\mathcal{C}, i) \subseteq P_i \cup Q_i$ and $\text{SPath}(\mathcal{C}, j) \subseteq P_j \cup Q_j$, we conclude that

$$\text{SPath}(\mathcal{C}, i) \cap \text{SPath}(\mathcal{C}, j) \subseteq (P_i \cap P_j) \cup (Q_i \cap Q_j).$$

We conclude the proof by noting that $P_i \cap P_j = \{p\}$ (\mathcal{P} is a simple geodesic tree) and $Q_i \cap Q_j = \text{SPath}(\mathcal{C}, i)[\text{lca}_{\mathcal{Q}}(a_i, a_j), \text{Root}(\mathcal{Q})]$. \square

In the example in Figure 20, paths $\text{Split}(\mathcal{C}, 2)$ and $\text{Split}(\mathcal{C}, 4)$ intersect at p and q only, whereas paths $\text{Split}(\mathcal{C}, 1)$ and $\text{Split}(\mathcal{C}, 4)$ intersect at p and their common suffix, $\text{Split}(\mathcal{C}, 4)[x, q]$ where $x = \text{lca}_{\mathcal{Q}}(a_1, a_4)$.

We will now introduce the concept of an order of a cage \mathcal{C} . Roughly speaking, if we let $\mathcal{C} = \text{Cage}((a_i)_{i=1}^{\ell}, (b_i)_{i=1}^{\ell}, \mathcal{P}, \mathcal{Q})$, then the paths $\text{Path}(\mathcal{P}, 1), \dots, \text{Path}(\mathcal{P}, \ell)$ originate at the common vertex $\text{Root}(\mathcal{P})$ and are vertex disjoint otherwise. Therefore, we can order all these paths cyclically around $\text{Root}(\mathcal{P})$. We stress that this cyclic order might not necessarily be consistent with the numbering of the paths. Fortunately, it turns out that we can order these paths in such a way that the paths $\text{Path}(\mathcal{Q}, 1), \dots, \text{Path}(\mathcal{Q}, \ell)$ are also sorted in a meaningful way. This is formalized by the following definition:

Definition 65. In a cage $\mathcal{C} = \text{Cage}((a_i)_{i=1}^{\ell}, (b_i)_{i=1}^{\ell}, \mathcal{P}, \mathcal{Q})$ of order ℓ in a graph G embedded in the plane, where $\mathcal{P} = \text{Tree}(P_1, \dots, P_{\ell})$, $\mathcal{Q} = \text{Tree}(Q_1, \dots, Q_{\ell})$, an **order** is a permutation $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_{\ell})$ of $\{1, 2, \dots, \ell\}$ such that:

- the first edges on paths $P_{\sigma_1}, P_{\sigma_2}, \dots, P_{\sigma_{\ell}}$ are ordered anti-clockwise around $\text{Root}(\mathcal{P})$ in the embedding of G in the plane, and
- for each vertex $v \in V(\mathcal{Q})$, there exist two indices L, R ($1 \leq L \leq R \leq \ell$) such that v lies on paths $Q_{\sigma_L}, Q_{\sigma_{L+1}}, \dots, Q_{\sigma_R}$ and on no other path Q_i .

For instance, the cage in Figure 20 has an order; permutations $\sigma = (1, 4, 5, 2, 3)$ and $\sigma = (5, 2, 3, 1, 4)$ are example orders. However, not all cyclic shifts of $(1, 4, 5, 2, 3)$ are correct orders — for instance, $\sigma = (4, 5, 2, 3, 1)$ is *not* an order because the vertex marked x in Figure 20 would belong to paths P_{σ_1} and P_{σ_5} and no other paths.

Using Definition 65, we can naturally refine the definition of a cage.

Definition 66. An **ordered cage** of order ℓ in a graph G embedded in the plane is a cage together with its order σ . We denote this structure as

$$\text{OrderedCage}((a_1, \dots, a_{\ell}), (b_1, \dots, b_{\ell}), \mathcal{P}, \mathcal{Q}, \sigma).$$

We are now ready to present a lemma which allows us to turn each cage into an ordered cage:

Lemma 67. *Each cage in a graph G embedded in the plane has an order.*

Proof. Fix a cage $\mathcal{C} = ((a_i)_{i=1}^{\ell}, (b_i)_{i=1}^{\ell}, \mathcal{P}, \mathcal{Q})$ in G . Define the objects:

$$\mathcal{P} = \text{Tree}(P_1, \dots, P_{\ell}), \quad \mathcal{Q} = \text{Tree}(Q_1, \dots, Q_{\ell}), \quad p = \text{Root}(\mathcal{P}), \quad q = \text{Root}(\mathcal{Q}).$$

If $\ell = 1$, the lemma is trivial. From now on, we assume that $\ell \geq 2$.

Consider a permutation $\mu^1 = (\mu_1^1, \mu_2^1, \dots, \mu_{\ell}^1)$ of $\{1, 2, \dots, \ell\}$ for which the first edges on paths $P_{\mu_1^1}, P_{\mu_2^1}, \dots, P_{\mu_{\ell}^1}$ are ordered anti-clockwise around p . We also consider all cyclic shifts of μ^1 — for each $k \in [1, \ell]$, we define $\mu^k = (\mu_1^k, \mu_2^k, \dots, \mu_{\ell}^k)$ as follows:

$$\mu^k = (\mu_k^1, \mu_{k+1}^1, \dots, \mu_{\ell}^1, \mu_1^1, \mu_2^1, \dots, \mu_{k-1}^1).$$

Clearly, for each $k \in [1, \ell]$, the first edges on the paths $P_{\mu_1^k}, P_{\mu_2^k}, \dots, P_{\mu_{\ell}^k}$ are ordered anti-clockwise around p .

We will prove that at least one of the permutations $\mu^1, \mu^2, \dots, \mu^{\ell}$ is an order. We begin by taking a vertex $v \in V(\mathcal{Q})$ and describing the set of paths Q_i containing v .

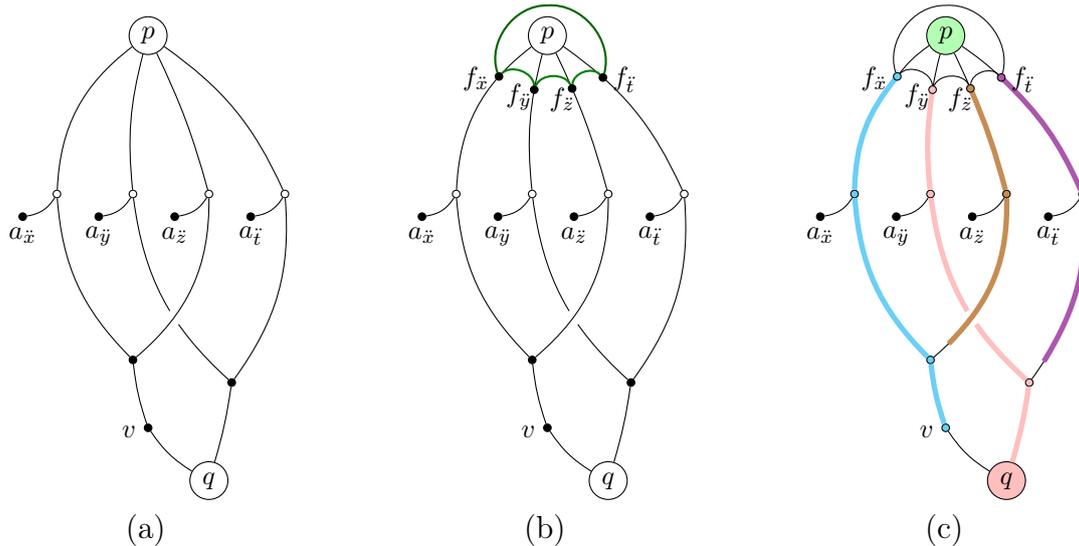


Figure 21: (a) — the setup in the proof of Claim 68.
 (b) — the edges added in G' (green).
 (c) — sets A_1, A_2, A_3, A_4, A_5 of vertices used in locating the K_5 clique as a minor of G' .

Claim 68. For every $v \in V(\mathcal{Q})$, no cyclic shift μ^k ($k \in [1, \ell]$) of μ^1 contains four distinct indices $\mu_x^k, \mu_y^k, \mu_z^k, \mu_t^k$ ($1 \leq x < y < z < t \leq \ell$) such that

$$v \in Q_{\mu_x^k}, \quad v \notin Q_{\mu_y^k}, \quad v \in Q_{\mu_z^k}, \quad v \notin Q_{\mu_t^k}.$$

Proof. For the sake of contradiction let us pick a vertex $v \in V(\mathcal{Q})$, an index $k \in [1, \ell]$, and indices $1 \leq x < y < z < t \leq \ell$ for which the conditions above hold. We simplify the notation by setting $\ddot{x} = \mu_x^k, \ddot{y} = \mu_y^k, \ddot{z} = \mu_z^k, \ddot{t} = \mu_t^k$ (Figure 21(a)).

We first remark that since q is the root of \mathcal{Q} , it belongs to all paths Q_i . Therefore, we must have that $v \neq q$.

We will create another planar graph G' which is a version of G with some edges added and some edges removed. Let $f_{\ddot{x}}, f_{\ddot{y}}, f_{\ddot{z}}, f_{\ddot{t}}$ be the second vertices on each of the paths $P_{\ddot{x}}, P_{\ddot{y}}, P_{\ddot{z}}$, and $P_{\ddot{t}}$, respectively — that is, the vertices immediately following p on the corresponding paths. We create a graph G' from G with the same set of vertices and the same embedding in the plane, but with the following modifications (Figure 21(b)):

- We disconnect p (the root of \mathcal{P}) from all its neighbors apart from $f_{\ddot{x}}, f_{\ddot{y}}, f_{\ddot{z}}, f_{\ddot{t}}$.
- We add the following edges to G' : $f_{\ddot{x}}f_{\ddot{y}}, f_{\ddot{y}}f_{\ddot{z}}, f_{\ddot{z}}f_{\ddot{t}}$, and $f_{\ddot{t}}f_{\ddot{x}}$.

Removing edges from the graph does not spoil its planarity. Afterwards, four edges are inserted to G' . In this graph, the anti-clockwise cyclic order of all edges incident to p is $pf_{\ddot{x}}, pf_{\ddot{y}}, pf_{\ddot{z}}, pf_{\ddot{t}}$. Each new edge $f_i f_j$ connects a neighbor of p with its next neighbor in this cyclic order. In order to preserve the planarity of G' , we draw an edge $f_i f_j$ so that the triangle created by edges $f_i f_j, f_j p, p f_i$ in this order is oriented anti-clockwise, and $f_i f_j$ lies close to edges $p f_i$ and $p f_j$ in the embedding (Figure 21(b)). Since p has only four neighbors in G' , we can draw each edge without breaking the planarity of the embedding of G' .

We will now find the clique K_5 as a minor of G' . We define five subgraphs of G' , induced by the following sets of vertices (Figure 21(c)):

$$\begin{aligned} A_1 &= \{p\}, \\ A_2 &= \text{SPath}(\mathcal{C}, \ddot{x})[f_{\ddot{x}}, v], \\ A_3 &= \text{SPath}(\mathcal{C}, \ddot{y})[f_{\ddot{y}}, q], \\ A_4 &= \text{SPath}(\mathcal{C}, \ddot{z})[f_{\ddot{z}}, v] - A_2, \\ A_5 &= \text{SPath}(\mathcal{C}, \ddot{t})[f_{\ddot{t}}, q] - A_3. \end{aligned}$$

Since $v \in Q_{\ddot{x}}$ and $v \in Q_{\ddot{z}}$, Lemma 63 asserts that $v \in \text{SPath}(\mathcal{C}, \ddot{x})$ and $v \in \text{SPath}(\mathcal{C}, \ddot{z})$. It is also obvious that all other ends of paths in the definitions of A_2, A_3, A_4, A_5 belong to the respective paths. Hence, the sets above are defined correctly. These sets are also non-empty since

$$p \in A_1, \quad f_{\ddot{x}} \in A_2, \quad f_{\ddot{y}} \in A_3, \quad f_{\ddot{z}} \in A_4, \quad f_{\ddot{t}} \in A_5.$$

We now prove that each of the induced subgraphs $G'[A_1], G'[A_2], G'[A_3], G'[A_4], G'[A_5]$ is connected. This fact is obvious for A_1, A_2 , and A_3 . For A_4 , we need to observe that because of Lemma 64, paths $\text{SPath}(\mathcal{C}, \ddot{x})$ and $\text{SPath}(\mathcal{C}, \ddot{z})$ intersect exactly at p and their common suffix. Therefore, removing A_2 from $\text{SPath}(\mathcal{C}, \ddot{z})[f_{\ddot{z}}, v]$ erases a suffix from the path, and leaves a subpath as a result. An analogous argument applies to A_5 . Therefore, each of the subgraphs is connected.

We can also prove the pairwise disjointness of A_1, \dots, A_5 . Obviously, A_1 is disjoint with all the remaining sets A_2, A_3, A_4, A_5 . Also, A_2 is disjoint with A_4 by the definition of A_4 , and A_3 is disjoint with A_5 by the definition of A_5 . Now, take any indices $i \in \{\ddot{x}, \ddot{z}\}$, $j \in \{\ddot{y}, \ddot{t}\}$, and let us prove that paths $\text{SPath}(\mathcal{C}, i)[f_i, v]$ and $\text{SPath}(\mathcal{C}, j)[f_j, q]$ are vertex-disjoint. By Lemma 64 we have that

$$\text{SPath}(\mathcal{C}, i) \cap \text{SPath}(\mathcal{C}, j) = \{p\} \cup \text{SPath}(\mathcal{C}, i)[\text{lca}_{\mathcal{Q}}(a_i, a_j), q] \quad (4)$$

(where $\text{SPath}(\mathcal{C}, i)[\text{lca}_{\mathcal{Q}}(a_i, a_j), q]$ is the common suffix of $\text{SPath}(\mathcal{C}, i)$ and $\text{SPath}(\mathcal{C}, j)$).

Naturally, $p \notin \text{SPath}(\mathcal{C}, i)[f_i, v]$. As $v \notin Q_j$, we get that $v \notin \text{SPath}(\mathcal{C}, i)[\text{lca}_{\mathcal{Q}}(a_i, a_j), q]$, so $\text{SPath}(\mathcal{C}, i)[f_i, v]$ is also disjoint with $\text{SPath}(\mathcal{C}, i)[\text{lca}_{\mathcal{Q}}(a_i, a_j), q]$. Hence, by (4), we infer that $\text{SPath}(\mathcal{C}, i)[f_i, v]$ is disjoint with $\text{SPath}(\mathcal{C}, j)$ as well. This means that we proved the vertex-disjointness of each pair of sets A_1, \dots, A_5 .

We are left to see that there is an edge between each pair of sets:

- There is a direct edge $pf_{\ddot{x}}$ connecting sets A_1 and A_2 ; we analogously show edges connecting A_1 with each remaining set.
- There is a direct edge $f_{\ddot{x}}f_{\ddot{y}}$ connecting sets A_2 and A_3 ; we analogously show edges connecting A_3 with A_4 , A_4 with A_5 , and A_5 with A_2 .
- Since the sets $A_2 = \text{SPath}(\mathcal{C}, \ddot{x})[f_{\ddot{x}}, v]$ and $\text{SPath}(\mathcal{C}, \ddot{z})[f_{\ddot{z}}, v]$ induce connected graphs and have a non-empty intersection, then naturally $A_4 = \text{SPath}(\mathcal{C}, \ddot{z})[f_{\ddot{z}}, v] - A_2$ is connected by an edge with A_2 . We analogously show that A_3 and A_5 are connected by an edge.

Summing everything up, we created a modified planar graph G' and indicated five non-empty connected subgraphs such that each pair of subgraphs is vertex-disjoint, yet there exist edges connecting each pair of subgraphs. Therefore, G' contains K_5 as a minor and hence — by Wagner's theorem [2] — it cannot be planar. The contradiction concludes the proof of the claim. \lrcorner

As a direct consequence of Claim 68, we can see that no cyclic shift μ^k of μ^1 can contain four direct indices $\mu_x^k, \mu_y^k, \mu_z^k, \mu_t^k$ (defined as in Claim 68) such that

$$v \notin Q_{\mu_x^k}, \quad v \in Q_{\mu_y^k}, \quad v \notin Q_{\mu_z^k}, \quad v \in Q_{\mu_t^k}.$$

If this set of conditions was true, we could consider the cyclic shift of μ^k with μ_y^k at the first position. We can easily see that this cyclic shift would be contradictory with Claim 68.

For a vertex $v \in V(\mathcal{Q})$ and integer $k \in [1, \ell]$, we set $\mathcal{K}_v^k := \{i \in [1, \ell] \mid v \in Q_{\mu_i^k}\}$; in other words, \mathcal{K}_v^k is the set of positions i in the permutation μ^k for which the corresponding path $Q_{\mu_i^k}$ contains v as a vertex.

Note that for every $k_1, k_2 \in [1, \ell]$ and every vertex $v \in V(\mathcal{Q})$, set $\mathcal{K}_v^{k_1}$ is a cyclic shift of $\mathcal{K}_v^{k_2}$; this is because the permutations μ^{k_1} and μ^{k_2} are also mutual cyclic shifts.

Claim 69. *Each set \mathcal{K}_v^k is a cyclic interval in $[1, \ell]$; that is, it is either an interval or a set of the form $[x, \ell] \cup [1, y]$ for some $x > y$.*

Proof. Fix a vertex $v \in V(\mathcal{Q})$, integer $k \in [1, \ell]$, and the corresponding set \mathcal{K}_v^k . Assume that \mathcal{K}_v^k is not a (non-cyclic) interval in $[1, \ell]$. In this setup, there must exist three indices x, y, z ($1 \leq x < y < z \leq \ell$) for which $x \in \mathcal{K}_v^k$, $y \notin \mathcal{K}_v^k$, and $z \in \mathcal{K}_v^k$. Without loss of generality, we can assume that $[x + 1, z - 1]$ is disjoint with \mathcal{K}_v^k .

Note that there cannot exist an index $t > z$ such that $t \notin \mathcal{K}_v^k$ as this would directly contradict Claim 68. Therefore, we have $[z, \ell] \subseteq \mathcal{K}_v^k$.

Also, there cannot exist an index $t < x$ such that $t \notin \mathcal{K}_v^k$ as this would pose a contradiction with the direct consequence of Claim 68. Thus, we have $[1, x] \subseteq \mathcal{K}_v^k$.

Since $[1, x] \subseteq \mathcal{K}_v^k$, $[x + 1, z - 1] \cap \mathcal{K}_v^k = \emptyset$, and $[z, \ell] \subseteq \mathcal{K}_v^k$, we have that $\mathcal{K}_v^k = [1, x] \cup [z, \ell]$. \lrcorner

In order to complete our proof, we also need to know how the cyclic intervals \mathcal{K}_v^k interact with each other. It turns out that these intervals form a laminar family:

Claim 70. *For any two vertices $u, v \in V(\mathcal{Q})$ and $k \in [1, \ell]$, the sets \mathcal{K}_u^k and \mathcal{K}_v^k are either disjoint, or one of them is a subset of the other.*

Proof. We remark that for $v \in V(\mathcal{Q})$, we have the following equivalent definition of \mathcal{K}_v^k :

$$\mathcal{K}_v^k = \{i \in [1, \ell] \mid a_{\mu_i^k} \text{ is in the subtree of } \mathcal{Q} \text{ rooted at } v\}.$$

This is true since the path $Q_{\mu_i^k}$ in the original definition of \mathcal{K}_v^k originates at $\text{Root}(\mathcal{Q})$ and terminates at $a_{\mu_i^k}$ (a leaf in \mathcal{Q}). Hence, v belongs to this path if and only if its endpoint — $a_{\mu_i^k}$ — is in the subtree rooted at v .

It is however well-known that the family of sets of leaves in every rooted tree T , where each set corresponds to a subset of leaves in a rooted subtree of T , is laminar. The proof of the claim follows. \lrcorner

Having proved all the necessary observations, we can find the cyclic shift μ^k of μ^1 which forms an order in the embedding of the graph. Fix any vertex $v^* \in V(\mathcal{Q})$ for which the cyclic interval $\mathcal{K}_{v^*}^1$ is not equal to $[1, \ell]$, but otherwise contains the largest number of indices. Such a vertex exists since $\ell \geq 2$, so the leaves in \mathcal{Q} do not cover the whole interval $[1, \ell]$.

We now take the cyclic shift μ^k for which $\mathcal{K}_{v^*}^k$ is a prefix of $[1, \ell]$; this is possible since $\mathcal{K}_{v^*}^1$ is a cyclic interval in μ^1 .

Claim 71. μ^k is an order in \mathcal{C} .

Proof. Let M be an integer such that $\mathcal{K}_{v^*}^k = [1, M]$. We will now prove that μ^k is an order. Firstly, since μ^k is a cyclic shift of μ^1 , the first edges on the paths $P_{\mu_1^k}, P_{\mu_2^k}, \dots, P_{\mu_\ell^k}$ form an anti-clockwise order around p .

In order to comply with Definition 65, we still need to prove that for each vertex $v \in V(\mathcal{Q})$, there exist two indices L, R ($1 \leq L \leq R \leq \ell$) such that v lies exactly on paths $Q_{\mu_L^k}, Q_{\mu_{L+1}^k}, \dots, Q_{\mu_R^k}$.

Fix a vertex $v \in \mathcal{Q}$. Note that \mathcal{K}_v^k is a cyclic interval (Claim 69). If $\mathcal{K}_v^k = [1, \ell]$, then the proposition we are proving is trivial. On the other hand, assume that $\mathcal{K}_v^k \neq [1, \ell]$. From our assumptions, we have that $|\mathcal{K}_v^1| \leq |\mathcal{K}_{v^*}^1|$, so also $|\mathcal{K}_v^k| \leq |\mathcal{K}_{v^*}^k|$. From Claim 70, we get that \mathcal{K}_v^k is either disjoint or fully contained within $\mathcal{K}_{v^*}^k$; it cannot contain $\mathcal{K}_{v^*}^k$ strictly due to $|\mathcal{K}_v^k| \leq |\mathcal{K}_{v^*}^k|$.

Recall that $\mathcal{K}_{v^*}^k = [1, M]$. If \mathcal{K}_v^k is disjoint with $\mathcal{K}_{v^*}^k$, then $\mathcal{K}_v^k \subseteq [M + 1, \ell]$, and \mathcal{K}_v^k is a non-cyclic interval. If \mathcal{K}_v^k is fully contained within $\mathcal{K}_{v^*}^k$, then $\mathcal{K}_v^k \subseteq [1, M]$, and again \mathcal{K}_v^k is a non-cyclic interval.

Hence, every $v \in V(\mathcal{Q})$ can be assigned a subinterval $\mathcal{K}_v^k = [L, R] \subseteq [1, \ell]$ such that v lies exactly on paths $Q_{\mu_L^k}, Q_{\mu_{L+1}^k}, \dots, Q_{\mu_R^k}$. This finishes the proof that μ^k is an order. \square

Claim 71 finishes the proof of the lemma. \square

While it is possible to proceed to the next parts of the proof using ordered cages, the fact that we need to use two different orderings — one imposed by the semi-ladder, and one imposed by the order σ in the ordered cage — will be excessively unwieldy. In order to remedy this inconvenience, we will introduce a variant of ordered cages unifying both orderings. This variant will lead us to a slightly weaker upper bound on the maximum semi-ladder size, though the bound will still remain asymptotically the same.

Definition 72. An **identity ordered cage** of order ℓ is an ordered cage

$$\text{OrderedCage}((a_1, \dots, a_\ell), (b_1, \dots, b_\ell), \mathcal{P}, \mathcal{Q}, \sigma),$$

where $\sigma = (1, 2, \dots, \ell)$ is the identity permutation. We denote this structure as

$$\text{IdOrderedCage}((a_1, \dots, a_\ell), (b_1, \dots, b_\ell), \mathcal{P}, \mathcal{Q}).$$

We immediately follow with a proof that large identity ordered cages exist in planar graphs given that large ordered cages exist. We, however, remark that the proof might alter the embedding of a graph in the plane.

Lemma 73. *If a graph G embedded in the plane contains an ordered cage of order $(\ell - 1)^2 + 1$ for $\ell \geq 1$, then there exists an embedding of G in the plane in which G contains an identity ordered cage of order ℓ .*

Proof. Take any graph G with an ordered cage \mathcal{C} of order $t := (\ell - 1)^2 + 1$:

$$\mathcal{C} = \text{OrderedCage}((a_i)_{i=1}^t, (b_i)_{i=1}^t, \mathcal{P}, \mathcal{Q}, \sigma).$$

By Erdős-Szekeres Theorem [40], σ contains either an increasing subsequence of length ℓ , or a decreasing subsequence of length ℓ .

If σ contains an increasing subsequence σ_{inc} of length ℓ , then a subset of \mathcal{C} with indices σ_{inc} is an identity ordered cage of order ℓ in G with the original embedding in the plane.

Assume now that σ contains a decreasing subsequence σ_{dec} of length ℓ . In this case, we modify the embedding of G in the plane by reflecting it through any line in the Euclidean plane. Since in the original embedding, the first edges on paths $P_{\sigma_1}, P_{\sigma_2}, \dots, P_{\sigma_t}$ were ordered anti-clockwise around $\text{Root}(\mathcal{P})$, these edges are ordered clockwise in the reflected embedding. Hence, the reversed permutation $\sigma^{\text{rev}} = (\sigma_t, \sigma_{t-1}, \dots, \sigma_1)$ is an order in the modified embedding, since the first edges on paths $P_{\sigma_t}, P_{\sigma_{t-1}}, \dots, P_{\sigma_1}$ are ordered anti-clockwise in the altered embedding. The remaining condition in Definition 66 can be verified in a straightforward way. Now, we notice that $\sigma_{\text{dec}}^{\text{rev}}$ (the reversed subsequence σ_{dec}) is an increasing subsequence in the newly formed order σ^{rev} . Hence, the subset of \mathcal{C} indicated by $\sigma_{\text{dec}}^{\text{rev}}$ in the transformed embedding is an identity ordered cage. \square

An example identity ordered cage is depicted in Figure 22(a).

In identity ordered cages, the following useful fact holds:

Lemma 74. *Fix an identity ordered cage $\text{IdOrderedCage}((a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, \mathcal{P}, \mathcal{Q})$ in a graph G and three indices i, j, k ($1 \leq i < j < k \leq \ell$). In the set $\{\text{lca}_{\mathcal{Q}}(a_i, a_j), \text{lca}_{\mathcal{Q}}(a_j, a_k)\}$, one of the elements is equal to $\text{lca}_{\mathcal{Q}}(a_i, a_k)$ and is an ancestor of the other element in the set (it can be, in particular, equal to it).*

Proof. Since both elements of the set $\{\text{lca}_{\mathcal{Q}}(a_i, a_j), \text{lca}_{\mathcal{Q}}(a_j, a_k)\}$ are the ancestors of a_j in \mathcal{Q} , one of these elements is an ancestor of the other. Without loss of generality, assume that $\text{lca}_{\mathcal{Q}}(a_i, a_j)$ is an ancestor of the other element of the set. This means that a_k is in the subtree rooted at $\text{lca}_{\mathcal{Q}}(a_i, a_j)$. As both a_i and a_k are in this subtree, we infer that $\text{lca}_{\mathcal{Q}}(a_i, a_j)$ is an ancestor of $\text{lca}_{\mathcal{Q}}(a_i, a_k)$.

Since $\text{lca}_{\mathcal{Q}}(a_i, a_k) \in Q_i$, $\text{lca}_{\mathcal{Q}}(a_i, a_k) \in Q_k$, and $j \in [i, k]$, by the definition of an identity ordered cage we also have $\text{lca}_{\mathcal{Q}}(a_i, a_k) \in Q_j$. Therefore, $\text{lca}_{\mathcal{Q}}(a_i, a_k)$ is also an ancestor of a_j in \mathcal{Q} . Hence, $\text{lca}_{\mathcal{Q}}(a_i, a_k)$ is an ancestor of $\text{lca}_{\mathcal{Q}}(a_i, a_j)$.

We conclude that $\text{lca}_{\mathcal{Q}}(a_i, a_j) = \text{lca}_{\mathcal{Q}}(a_i, a_k)$. \square

5.3 Neighbor areas

Until now, we were concerned about the structure imposed only by some vertices of the semi-ladder — that is, the vertices a_1, a_2, \dots, a_ℓ . It is now time to include the corresponding vertices b_1, b_2, \dots, b_ℓ in the picture. It turns out that each sufficiently large ordered cage contains a smaller ordered cage as a subset in which for every i , the vertex b_i is, in some topological sense, *near* a_i . This is where the upper bounds on the neighborhood complexity in planar graphs proved in Section 4, will come into play.

First, we need to design technical tools allowing us to formalize this idea.

Definition 75. Let $\mathcal{C} = \text{IdOrderedCage}((a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, \mathcal{P}, \mathcal{Q})$ be an identity ordered cage of order ℓ in G . For every two integers i, j ($1 \leq i < j \leq \ell$), we set $v_{i,j} = \text{lca}_{\mathcal{Q}}(a_i, a_j)$. Note that by Lemma 64, $v_{i,j}$ is the first vertex of the common suffix of $\text{SPath}(\mathcal{C}, i)$ and $\text{SPath}(\mathcal{C}, j)$.

For these integers i, j , we define the **neighbor area** $\text{Area}(\mathcal{C}, i, j)$ as the closed part of the plane whose boundary is traversed anti-clockwise by the following directed cycle:

$$\text{SPath}(\mathcal{C}, i)[\text{Root}(\mathcal{P}), v_{i,j}] \cdot (\text{SPath}(\mathcal{C}, j))^{-1}[v_{i,j}, \text{Root}(\mathcal{P})].$$

Note that neighbor areas are not necessarily bounded (Figure 22).

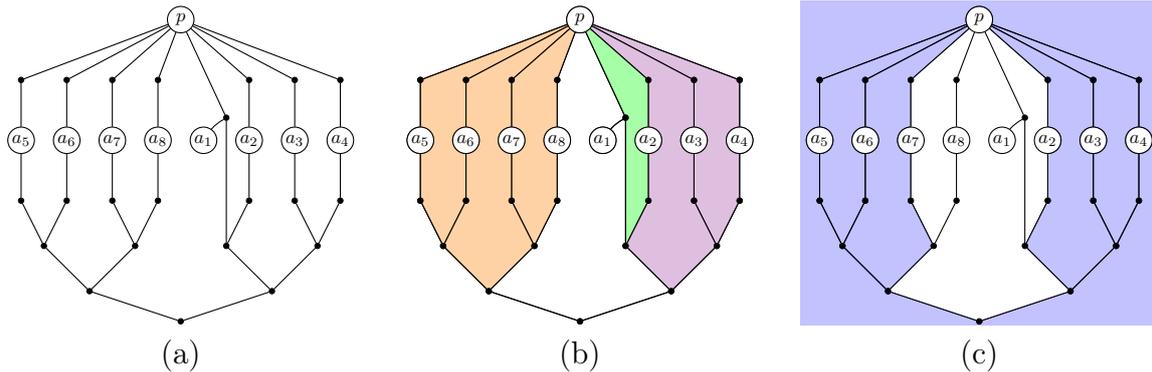


Figure 22: (a) — an example identity ordered cage \mathcal{C} of order 8 embedded in the plane. (b) — neighbor areas $\text{Area}(\mathcal{C}, 1, 2)$ (green), $\text{Area}(\mathcal{C}, 2, 4)$ (violet), and $\text{Area}(\mathcal{C}, 5, 8)$ (orange). (c) — an unbounded neighbor area $\text{Area}(\mathcal{C}, 2, 7)$ (blue). The area is unbounded since its boundary is traversed anti-clockwise first by $\text{SPath}(\mathcal{C}, 2)$, and later by $(\text{SPath}(\mathcal{C}, 7))^{-1}$.

We can easily see from Figure 22 that the union of two neighbor areas $\text{Area}(\mathcal{C}, 1, 2)$ and $\text{Area}(\mathcal{C}, 2, 4)$ in the presented example is $\text{Area}(\mathcal{C}, 1, 4)$, and that the intersection of these is a prefix of $\text{SPath}(\mathcal{C}, 2)$. It turns out that this observation is satisfied in general:

Lemma 76. *In an identity ordered cage \mathcal{C} of order ℓ , for every indices i, j, k such that $1 \leq i < j < k \leq \ell$, we have that:*

$$\begin{aligned} \text{Area}(\mathcal{C}, i, j) \cup \text{Area}(\mathcal{C}, j, k) &= \text{Area}(\mathcal{C}, i, k), \\ \text{Area}(\mathcal{C}, i, j) \cap \text{Area}(\mathcal{C}, j, k) &= \partial\text{Area}(\mathcal{C}, i, j) \cap \partial\text{Area}(\mathcal{C}, j, k) \subseteq \text{SPath}(\mathcal{C}, j). \end{aligned}$$

Proof. We first define vertices $v_{i,j} = \text{lca}_{\mathcal{Q}}(a_i, a_j)$, $v_{j,k} = \text{lca}_{\mathcal{Q}}(a_j, a_k)$ as in the definition of neighbor areas. We also let $\mathcal{C} = \text{IdOrderedCage}((a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, \mathcal{P}, \mathcal{Q})$, $p := \text{Root}(\mathcal{P})$, and $\mathcal{Q} = \text{Tree}(Q_1, \dots, Q_\ell)$. These definitions tell us that

$$\begin{aligned} \partial\text{Area}(\mathcal{C}, i, j) &= \text{SPath}(\mathcal{C}, i)[p, v_{i,j}] \cdot (\text{SPath}(\mathcal{C}, j))^{-1}[v_{i,j}, p], \\ \partial\text{Area}(\mathcal{C}, j, k) &= \text{SPath}(\mathcal{C}, j)[p, v_{j,k}] \cdot (\text{SPath}(\mathcal{C}, k))^{-1}[v_{j,k}, p]. \end{aligned}$$

Since $v_{i,j}$ and $v_{j,k}$ are both ancestors of a_j in \mathcal{Q} , either of these is an ancestor of the other (in particular, we might have $v_{i,j} = v_{j,k}$). Without loss of generality, we assume that $v_{j,k}$ is an ancestor of $v_{i,j}$ in \mathcal{Q} or equal to $v_{i,j}$ (Figure 23). Then, $\text{SPath}(\mathcal{C}, j)[p, v_{i,j}] \subseteq \text{SPath}(\mathcal{C}, j)[p, v_{j,k}]$. Moreover, by Lemma 74, $v_{i,k} := \text{lca}(a_i, a_k)$ is equal to $v_{j,k}$.

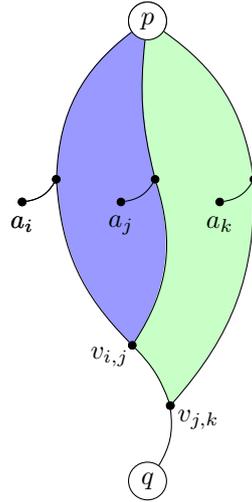


Figure 23: The configuration assumed in Lemma 76. The vertices $v_{i,j}$ and $v_{j,k}$ may overlap. The blue area is $\text{Area}(\mathcal{C}, i, j)$. The green area is $\text{Area}(\mathcal{C}, j, k)$. The brown path is $\text{SPath}(\mathcal{C}, j)$.

We first note that $\partial\text{Area}(\mathcal{C}, i, j) \cap \partial\text{Area}(\mathcal{C}, j, k) \subseteq \text{SPath}(\mathcal{C}, j)$; this is because any intersection of both boundaries outside of $\text{SPath}(\mathcal{C}, j)$ would have to be an intersection of $\text{SPath}(\mathcal{C}, i)$ with $\text{SPath}(\mathcal{C}, k)$. These paths intersect only at p and $\text{SPath}(\mathcal{C}, i)[v_{i,k}, q]$ (Lemma 64). However, $v_{i,k} = v_{j,k} \in \text{SPath}(\mathcal{C}, j)$, so also $\text{SPath}(\mathcal{C}, i)[v_{i,k}, q] \subseteq \text{SPath}(\mathcal{C}, j)$. Obviously, $p \in \text{SPath}(\mathcal{C}, j)$ as well. Therefore, no intersection of $\partial\text{Area}(\mathcal{C}, i, j)$ and $\partial\text{Area}(\mathcal{C}, j, k)$ can occur outside of $\text{SPath}(\mathcal{C}, j)$.

Thanks to this fact, we can now compute $\partial\text{Area}(\mathcal{C}, i, j) \cap \partial\text{Area}(\mathcal{C}, j, k)$:

$$\begin{aligned} \partial\text{Area}(\mathcal{C}, i, j) \cap \partial\text{Area}(\mathcal{C}, j, k) &= \\ &= (\partial\text{Area}(\mathcal{C}, i, j) \cap \text{SPath}(\mathcal{C}, j)) \cap (\partial\text{Area}(\mathcal{C}, j, k) \cap \text{SPath}(\mathcal{C}, j)) = \\ &= \text{SPath}(\mathcal{C}, j)[p, v_{i,j}] \cap \text{SPath}(\mathcal{C}, j)[p, v_{j,k}] = \text{SPath}(\mathcal{C}, j)[p, v_{i,j}] \end{aligned}$$

— that is, this intersection is a prefix of $\text{SPath}(\mathcal{C}, j)$.

Therefore, the two boundaries intersect at exactly one segment. We can use this information to prove that the interiors of $\text{Area}(\mathcal{C}, i, j)$ and $\text{Area}(\mathcal{C}, j, k)$ are disjoint. Let us name f_i, f_j, f_k as the second vertices on the paths P_i, P_j, P_k , respectively; that is, the vertices immediately following $p = \text{Root}(\mathcal{P})$ on the respective paths. Since the counter-clockwise orientation of $\partial\text{Area}(\mathcal{C}, i, j)$ contains edges $f_j p$ and $p f_i$ in this order, and the counter-clockwise orientation of $\partial\text{Area}(\mathcal{C}, j, k)$ contains edges $f_k p$ and $p f_j$ in this order, we only need to verify the cyclic order of the edges $p f_i, p f_j, p f_k$ around vertex p . These edges are ordered anti-clockwise around p (by the definition of an ordered cage), so the interiors of the areas $\text{Area}(\mathcal{C}, i, j)$, $\text{Area}(\mathcal{C}, j, k)$ are disjoint (Figure 23).

From the considerations above, we conclude that

$$\text{Area}(\mathcal{C}, i, j) \cap \text{Area}(\mathcal{C}, j, k) = \partial\text{Area}(\mathcal{C}, i, j) \cap \partial\text{Area}(\mathcal{C}, j, k) \subseteq \text{SPath}(\mathcal{C}, j).$$

Moreover, since the interiors of both neighbor areas are disjoint, given the anti-clockwise orientations of $\partial\text{Area}(\mathcal{C}, i, j)$ and $\partial\text{Area}(\mathcal{C}, j, k)$, both starting and finishing at

their common vertex p , we can compute the anti-clockwise orientation of

$$\partial[\text{Area}(\mathcal{C}, i, j) \cup \text{Area}(\mathcal{C}, j, k)]$$

by composing the anti-clockwise orientations of $\partial\text{Area}(\mathcal{C}, i, j)$ and $\partial\text{Area}(\mathcal{C}, j, k)$:

$$\begin{aligned} \partial[\text{Area}(\mathcal{C}, i, j) \cup \text{Area}(\mathcal{C}, j, k)] &= \\ &= \text{SPath}(\mathcal{C}, i)[p, v_{i,j}] \cdot \text{SPath}(\mathcal{C}, j)^{-1}[v_{i,j}, p] \cdot \text{SPath}(\mathcal{C}, j)[p, v_{j,k}] \cdot \text{SPath}(\mathcal{C}, k)^{-1}[v_{j,k}, p] = \\ &= \text{SPath}(\mathcal{C}, i)[p, v_{i,j}] \cdot \text{SPath}(\mathcal{C}, j)[v_{i,j}, v_{j,k}] \cdot \text{SPath}(\mathcal{C}, k)^{-1}[v_{j,k}, p] \stackrel{(\star)}{=} \\ &\stackrel{(\star)}{=} \text{SPath}(\mathcal{C}, i)[p, v_{i,j}] \cdot \text{SPath}(\mathcal{C}, i)[v_{i,j}, v_{j,k}] \cdot \text{SPath}(\mathcal{C}, k)^{-1}[v_{j,k}, p] = \\ &= \text{SPath}(\mathcal{C}, i)[p, v_{j,k}] \cdot \text{SPath}(\mathcal{C}, k)^{-1}[v_{j,k}, p] = \\ &= \text{SPath}(\mathcal{C}, i)[p, v_{i,k}] \cdot \text{SPath}(\mathcal{C}, k)^{-1}[v_{i,k}, p] = \partial\text{Area}(\mathcal{C}, i, k). \end{aligned}$$

In the equality (\star) , we use the definition of $v_{i,j}$ — it is the first vertex of the common suffix of $\text{SPath}(\mathcal{C}, i)$ and $\text{SPath}(\mathcal{C}, j)$.

Therefore, $\text{Area}(\mathcal{C}, i, j) \cup \text{Area}(\mathcal{C}, j, k) = \text{Area}(\mathcal{C}, i, k)$. □

We can immediately infer the following fact from Lemma 76:

Corollary 77. *In any identity ordered cage \mathcal{C} of order ℓ , for every triple of indices i, j, k such that $1 \leq i < j < k \leq \ell$, we have that:*

$$\text{Area}(\mathcal{C}, i, j) \cap \text{Int Area}(\mathcal{C}, j, k) = \emptyset, \quad \text{Int Area}(\mathcal{C}, i, j) \cap \text{Area}(\mathcal{C}, j, k) = \emptyset.$$

Proof. Since areas $\text{Area}(\mathcal{C}, i, j)$ and $\text{Area}(\mathcal{C}, j, k)$ intersect only at their boundaries, the interior of either neighbor area is disjoint with the other neighbor area. □

Now we will examine the importance of neighbor areas in identity ordered cages. For an identity ordered cage \mathcal{C} of order ℓ and two indices i, j ($i < j$), we can observe which vertices of the semi-ladder underlying \mathcal{C} : $a_1, a_2, \dots, a_\ell, b_1, b_2, \dots, b_\ell$ belong to $\text{Area}(\mathcal{C}, i, j)$, and which do not. It turns out that there cannot be too many pairs of vertices (a_k, b_k) for $k \in [1, \ell]$ such that one vertex in the pair belongs to $\text{Area}(\mathcal{C}, i, j)$ and the other does not. For simplicity of the further part of the proof, we use the following definition:

Definition 78. Fix a neighbor area $\text{Area}(\mathcal{C}, i, j)$ in an identity ordered cage \mathcal{C} . For a pair of vertices u, v , if exactly one of them belongs to $\text{Area}(\mathcal{C}, i, j)$, we will say that u is **separated from** v by $\text{Area}(\mathcal{C}, i, j)$.

For instance, in Figure 22, a_5 is separated from a_3 by $\text{Area}(\mathcal{C}, 2, 4)$, but a_5 is not separated from a_7 by the same neighbor area.

It turns out that the definition of “being separated” is connected in a non-trivial way to the distinctness of distance- d profiles. This will allow us to employ the Noose Profile Lemma (Theorem 28), which bounds the neighborhood complexity of nooses in planar graphs, in order to limit the number of pairs of the vertices of the semi-ladder separated by neighbor areas.

Lemma 79. *In an identity ordered cage $\mathcal{C} = \text{IdOrderedCage}((a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, \mathcal{P}, \mathcal{Q})$ of order ℓ , for every pair of indices i, j ($1 \leq i < j \leq \ell$), the neighbor area $\text{Area}(\mathcal{C}, i, j)$ separates at most $128d^3(d+2)^4$ distinct pairs of vertices (a_k, b_k) for $k \in [1, \ell]$.*

Proof. Fix two integers i, j ($1 \leq i < j \leq \ell$) for the proof. Denote the set of vertices on the boundary of $\text{Area}(\mathcal{C}, i, j)$ as B .

Claim 80. $|B| \leq 4d$.

Proof. We first note that $|E(\text{SPath}(\mathcal{C}, k))| \leq 2d$ for each $k \in [1, L]$ since $\text{SPath}(\mathcal{C}, k)$ is the concatenation of subpaths of two paths, $\text{Path}(\mathcal{P}, k)$ and $\text{Path}(\mathcal{Q}, k)$, whose lengths do not exceed d . The boundary of $\text{Area}(\mathcal{C}, i, j)$ is a cycle and consists of two prefixes of paths $\text{SPath}(\mathcal{C}, i)$ and $\text{SPath}(\mathcal{C}, j)$. Therefore,

$$|B| = |V(\partial \text{Area}(\mathcal{C}, i, j))| = |E(\partial \text{Area}(\mathcal{C}, i, j))| \leq |E(\text{SPath}(\mathcal{C}, i))| + |E(\text{SPath}(\mathcal{C}, j))| \leq 4d.$$

□

Let I denote the set of indices k for which a_k is separated from b_k by $\text{Area}(\mathcal{C}, i, j)$; we need to prove that $|I| \leq 128d^3(d+2)^4$.

For each $k \in I$, two cases are possible — either a_k belongs to $\text{Area}(\mathcal{C}, i, j)$, which implies that $b_k \notin \text{Area}(\mathcal{C}, i, j)$; or $a_k \notin \text{Area}(\mathcal{C}, i, j)$, which implies that $b_k \in \text{Area}(\mathcal{C}, i, j)$. We will partition I into two subsets — I_1 and I_2 — depending on which case holds for a given k :

$$I_1 = \{k \in I \mid a_k \in \text{Area}(\mathcal{C}, i, j), b_k \notin \text{Area}(\mathcal{C}, i, j)\},$$

$$I_2 = \{k \in I \mid a_k \notin \text{Area}(\mathcal{C}, i, j), b_k \in \text{Area}(\mathcal{C}, i, j)\}.$$

We will prove that both I_1 and I_2 contain no more than $64d^3(d+2)^4$ indices; the statement of the lemma will follow immediately.

Claim 81. For any $t \in \{1, 2\}$ and two different indices $x, y \in I_t$, vertices a_x and a_y have different distance- d profiles on B .

Proof. Assume on the contrary that a_x and a_y have identical distance- d profiles on B . Without loss of generality, assume that $x < y$. Because vertices a_x, a_y, b_x, b_y form a distance- d semi-ladder of order 2, it follows that

$$\text{dist}(a_y, b_x) \leq d, \quad \text{but} \quad \text{dist}(a_x, b_x) > d.$$

Since $x, y \in I_k$, we have that a_y is separated from b_x by $\text{Area}(\mathcal{C}, i, j)$. Therefore, the

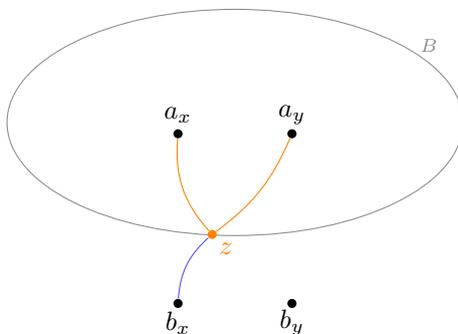


Figure 24: The configuration in Claim 81.

shortest path between these two vertices must pass through a vertex of $\partial\text{Area}(\mathcal{C}, i, j)$ — that is, some vertex $z \in B$ (Figure 24). Thus,

$$\text{dist}(a_y, b_x) = \text{dist}(a_y, z) + \text{dist}(z, b_x).$$

It follows that $\text{dist}(a_y, z) \leq d$ and hence $\pi_d[a_y, B](z) = \text{dist}(a_y, z)$. Because distance- d profiles of a_y and a_x on B are equal, this implies that $\pi_d[a_x, B](z) = \pi_d[a_y, B](z) = \text{dist}(a_y, z)$. Since also $\pi_d[a_x, B](z) \leq d$, we get that $\text{dist}(a_x, z) = \pi_d[a_x, B](z) = \text{dist}(a_y, z)$, so the following inequality follows from the triangle inequality:

$$\text{dist}(a_x, b_x) \leq \text{dist}(a_x, z) + \text{dist}(z, b_x) = \text{dist}(a_y, z) + \text{dist}(z, b_x) = \text{dist}(a_y, b_x) \leq d$$

— a contradiction. Therefore, a_x and a_y must have distinct distance- d profiles on B . \lrcorner

A combination of both claims above provides us with a link between the notion of cages and the Noose Profile Lemma. Assume for contradiction that for an integer $t \in \{1, 2\}$, the set I_t contains more than $64d^3(d+2)^4$ elements. By Claim 81, we infer that $|\{\pi_d[a_x, B] \mid x \in I_k\}| > 64d^3(d+2)^4$. We also remark that since B is a cycle in G , there exists a noose \mathcal{L} in the embedding of G in the plane passing through vertices in B and no other vertices (\mathcal{L} should closely follow the cycle B , passing through each vertex of B in the same order as B , but not touching the edges of the graph). \mathcal{L} is now subject to the Noose Profile Lemma. We can see that either all vertices in $\{a_k \mid k \in I_t\}$ are enclosed by \mathcal{L} , or none of them are enclosed by \mathcal{L} (depending on the value of t and whether $\text{Area}(\mathcal{C}, i, j)$ is bounded or unbounded). If these vertices are all enclosed by \mathcal{L} , then Noose Profile Lemma (Theorem 28) applies, so the number of different distance- d profiles seen from \mathcal{L} is bounded from above by

$$|B|^3 \cdot (d+2)^4 \leq (4d)^3 \cdot (d+2)^4 = 64d^3(d+2)^4$$

— a contradiction. Analogously, if none of vertices in $\{a_k \mid k \in I_t\}$ are enclosed by \mathcal{L} , then a variant of the Noose Profile Lemma (Corollary 49) applies, and we contradict our assumption in the same fashion. Therefore, $|I_t| \leq 64d^3(d+2)^4$.

Hence, $|I| = |I_1| + |I_2| \leq 128d^3(d+2)^4$. \square

We will also prove that in an identity ordered cage \mathcal{C} , for each $i \in [2, \ell - 1]$, we can expect to find the vertex a_i in the interior of the neighbor area $\text{Area}(\mathcal{C}, i - 1, i + 1)$ — roughly speaking, “near” the path $\text{SPath}(\mathcal{C}, i)$.

Lemma 82. *In an identity ordered cage*

$$\mathcal{C} = \text{IdOrderedCage}((a_1, \dots, a_\ell), (b_1, \dots, b_\ell), \mathcal{P}, \mathcal{Q})$$

in graph G , for every $i \in [2, \ell - 1]$, we have that $a_i \in \text{Int Area}(\mathcal{C}, i - 1, i + 1)$.

Proof. For the sake of contradiction assume that $a_i \notin \text{Int Area}(\mathcal{C}, i - 1, i + 1)$.

For all $j \in [1, \ell]$, we choose f_j as the second vertex on $\text{Path}(\mathcal{P}, j)$ — that is, the vertex on this path immediately following $\text{Root}(\mathcal{P})$. We also let $p := \text{Root}(\mathcal{P})$, $\mathcal{P} = \text{Tree}(P_1, \dots, P_\ell)$, and $\mathcal{Q} = \text{Tree}(Q_1, \dots, Q_\ell)$.

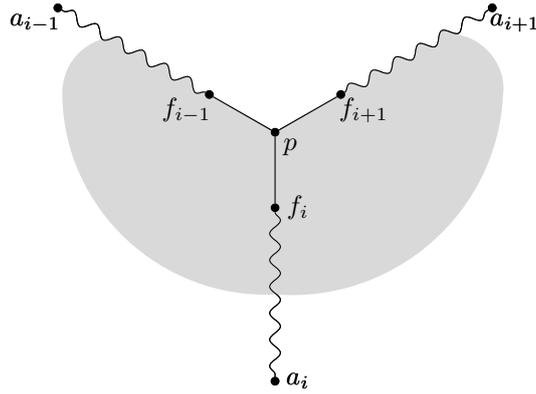


Figure 25: A simple geodesic tree \mathcal{P} (a part of the identity ordered cage \mathcal{C}) with important vertices in Lemma 82. $\text{Area}(\mathcal{C}, i - 1, i + 1)$ is shaded.

Since edges pf_{i-1}, pf_i, pf_{i+1} are ordered anti-clockwise around p (by the definition of an ordered cage), and because the boundary of $\text{Area}(\mathcal{C}, i - 1, i + 1)$, when ordered anti-clockwise, contains edges $f_{i+1}p$ and pf_{i-1} in this order, we have that $pf_i \subseteq \text{Area}(\mathcal{C}, i - 1, i + 1)$ (Figure 25). Thus, $f_i \in \text{Area}(\mathcal{C}, i - 1, i + 1)$. Since $a_i \notin \text{Int Area}(\mathcal{C}, i - 1, i + 1)$, the path $P_i[f_i, a_i]$ must intersect $\partial \text{Area}(\mathcal{C}, i - 1, i + 1)$.

By the definition of neighbor areas, $\partial \text{Area}(\mathcal{C}, i - 1, i + 1) \subseteq P_{i-1} \cup P_{i+1} \cup Q_{i-1} \cup Q_{i+1}$. However, $P_i[f_i, a_i]$ is vertex-disjoint with P_j for each $j \neq i$ (since by the definition of a quasi-cage, paths P_i and P_j are vertex-disjoint apart from p). Similarly, P_i is vertex-disjoint with Q_j for each $j \neq i$ (by the definition of a cage). Hence, $P_i[f_i, a_i]$ is vertex-disjoint with $\partial \text{Area}(\mathcal{C}, i - 1, i + 1)$ — a contradiction. \square

5.4 Neighbor cages

We are now ready to use the tools developed in Section 5.3 to define neighbor cages and explore their properties.

Definition 83. An identity ordered cage $\mathcal{C} = \text{IdOrderedCage}((a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, \mathcal{P}, \mathcal{Q})$ of order ℓ is also a **neighbor cage** if for each index $i \in \{2, 3, \dots, \ell - 1\}$, the vertex b_i belongs to $\text{Area}(\mathcal{C}, i - 1, i + 1)$.

We denote neighbor cages in the following way:

$$\text{NeighborCage}((a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, \mathcal{P}, \mathcal{Q}).$$

Figure 26 shows an example of a neighbor cage. Unnecessary parts of the graph (including the edges connecting the vertices b_i with the remaining part of the graph) have been omitted.

We will now employ the toolchain we have developed to prove that large identity ordered cages contain large neighbor cages as subsets.

Lemma 84. *Each identity ordered cage of order $[256d^3(d + 2)^4 + 2]\ell + 2$ contains a neighbor cage of order ℓ as a subset.*

Proof. Let $L := [256d^3(d + 2)^4 + 1]\ell + 2$, and fix an identity ordered cage of order L :

$$\mathcal{C} = \text{IdOrderedCage}((a_i)_{i=1}^L, (b_i)_{i=1}^L, \mathcal{P}, \mathcal{Q}).$$

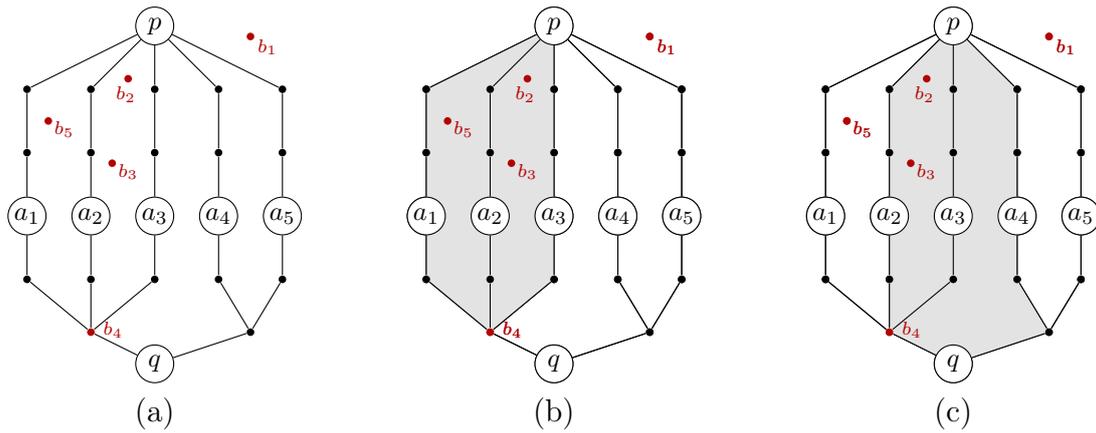


Figure 26: (a) — An example neighbor cage of order 5, with edges connecting b_1, \dots, b_5 with the remaining part of the graph removed. Note that b_1 and b_5 may be anywhere in the graph.

(b) — The neighbor area which must contain b_2 is shaded.

(c) — The neighbor area which must contain b_3 is shaded.

We build a permutation π of vertices of the semi-ladder $\{a_1, a_2, \dots, a_L, b_1, b_2, \dots, b_L\}$ in the following way:

- π contains vertices a_1, a_2, \dots, a_L in this order; immediately after each a_i (for $i \in [1, L]$), we add a (possibly empty) sequence Γ_i which is a permuted subset of $\{b_1, b_2, \dots, b_L\}$. That is,

$$\pi = (a_1) \cdot \Gamma_1 \cdot (a_2) \cdot \Gamma_2 \cdots \cdots (a_L) \cdot \Gamma_L,$$

where \cdot denotes the concatenation of sequences.

- For each index $i \in [1, L - 1]$, Γ_i contains those vertices b_1, \dots, b_L which belong to $\text{Area}(\mathcal{C}, i, i + 1)$, but do not belong to any previous sequence $\Gamma_1, \dots, \Gamma_{i-1}$. Formally, the set of vertices in Γ_i is

$$[\{b_1, b_2, \dots, b_L\} \cap \text{Area}(\mathcal{C}, i, i + 1)] - (\Gamma_1 \cup \dots \cup \Gamma_{i-1}).$$

- Γ_L contains all vertices in $\{b_1, b_2, \dots, b_L\}$ which do not belong to any previous sequence $\Gamma_1, \dots, \Gamma_{L-1}$.
- Each sequence Γ_i (for $i \in [1, L]$) is ordered arbitrarily.

It can be clearly seen that the procedure above produces a permutation π .

We will now find a large subsequence of π with the following property: if a_i belongs to the subsequence for some $i \in [1, L]$, then b_i is a neighbor of a_i in this subsequence; and conversely, if b_i belongs to this subsequence, then a_i neighbors b_i in this subsequence.

For all integers $i \in [1, L]$, we define the values $\alpha(i)$ and $\beta(i)$ so that $a_i = \pi_{\alpha(i)}$ and $b_i = \pi_{\beta(i)}$. Note that by the definition of π , we have $\alpha(1) < \alpha(2) < \dots < \alpha(L)$.

This definition allows us to define the **conflict interval** X_i for each $i \in [2, L - 1]$ — a closed interval whose endpoints are at $\alpha(i)$ and $\beta(i)$. Each conflict interval is a subinterval of the interval $[1, 2L]$. We stress that we do not define X_1 or X_L .

We note the following property of the conflict intervals:

Claim 85. *For every two different integers s, t ($s \in [2, L - 1], t \in [2, L]$), a_s is separated from b_s by $\text{Area}(\mathcal{C}, 1, t)$ if and only if $\alpha(t) \in X_s$.*

Proof. Let us first remark that

$$\text{Area}(\mathcal{C}, 1, t) = \text{Area}(\mathcal{C}, 1, 2) \cup \text{Area}(\mathcal{C}, 2, 3) \cup \dots \cup \text{Area}(\mathcal{C}, t - 1, t) \quad (5)$$

(Lemma 76). Therefore, if $\beta(s) < \alpha(t)$, then b_s was added to π before a_t ; it follows that b_s belongs to one of the subsequences $\Gamma_1, \Gamma_2, \dots, \Gamma_{t-1}$. Hence, b_s belongs to $\text{Area}(\mathcal{C}, i, i + 1)$ for some $i < t$. By (5), we conclude that $b_s \in \text{Area}(\mathcal{C}, 1, t)$. Conversely, if $b_s \in \text{Area}(\mathcal{C}, 1, t)$, then $b_s \in \text{Area}(\mathcal{C}, i, i + 1)$ for some $i < t$. If we fix i to be the smallest index for which $b_s \in \text{Area}(\mathcal{C}, i, i + 1)$, we will get that $b_s \in \Gamma_i$, where $i < t$, so also $\beta(s) < \alpha(t)$. We conclude that

$$b_s \in \text{Area}(\mathcal{C}, 1, t) \Leftrightarrow \beta(s) < \alpha(t). \quad (6)$$

Now, let us compare $\alpha(s)$ with $\alpha(t)$. Naturally, we have $\alpha(s) < \alpha(t)$ if and only if $s < t$. Also, by Lemma 82 we have that $a_s \in \text{Int Area}(\mathcal{C}, s - 1, s + 1)$. If $s < t$, we immediately get that $a_s \in \text{Area}(\mathcal{C}, 1, t)$. However, if $s > t$, then we infer the following relation from Corollary 77:

$$\text{Area}(\mathcal{C}, 1, s - 1) \cap \text{Int Area}(\mathcal{C}, s - 1, s + 1) = \emptyset.$$

Hence, $a_s \notin \text{Area}(\mathcal{C}, 1, s - 1)$, so also $a_s \notin \text{Area}(\mathcal{C}, 1, t)$. We now proved that

$$a_s \in \text{Area}(\mathcal{C}, 1, t) \Leftrightarrow \alpha(s) < \alpha(t). \quad (7)$$

It is obvious that $\alpha(t) \in X_s$ if and only if exactly one of the conditions $\alpha(s) < \alpha(t)$, $\beta(s) < \alpha(t)$ holds. From (6) and (7) it follows that this is equivalent to exactly one of the conditions $a_s \in \text{Area}(\mathcal{C}, 1, t)$, $b_s \in \text{Area}(\mathcal{C}, 1, t)$ being satisfied. A verification with Definition 78 concludes the proof. \lrcorner

This claim enables us to prove that the conflict intervals do not intersect *too frequently*:

Claim 86. *For each $x \in \mathbb{R}$, at most $256d^3(d + 2)^4 + 2$ conflict intervals contain x .*

Proof. Fix $x \in \mathbb{R}$, and consider the set $I = \{i \in [2, L - 1] \mid x \in X_i\}$ of all indices for which the corresponding conflict interval contains x . We partition I into two subsets, depending on whether $\alpha(i)$ is the left or the right endpoint of the conflict interval X_i :

$$I_{\alpha\beta} = \{i \in I \mid \alpha(i) < \beta(i)\},$$

$$I_{\beta\alpha} = \{i \in I \mid \alpha(i) > \beta(i)\}.$$

Let $m = |I_{\alpha\beta}|$ and $I_{\alpha\beta} = \{i_1, i_2, \dots, i_m\}$ where $2 \leq i_1 < i_2 < \dots < i_m \leq L - 1$. Since $\alpha(i_1) < \alpha(i_2) < \dots < \alpha(i_m)$ and each interval X_{i_j} for $j \in [1, m]$ has its left endpoint at $\alpha(i_j)$, we conclude that $\alpha(i_m)$ also belongs to each conflict interval $X_{i_1}, X_{i_2}, \dots, X_{i_m}$. Hence, by Claim 85, each pair of vertices a_{i_j}, b_{i_j} for $j \in [1, m - 1]$ is separated by

$\text{Area}(\mathcal{C}, 1, i_m)$. By Lemma 79 we conclude that $m - 1 \leq 128d^3(d + 2)^4$, so $|I_{\alpha\beta}| \leq 128d^3(d + 2)^4 + 1$.

We can prove in a similar fashion that $|I_{\beta\alpha}| \leq 128d^3(d + 2)^4 + 1$: if $I_{\beta\alpha} = \{j_1, j_2, \dots, j_n\}$ for some $n \in \mathbb{N}$ and $2 \leq j_1 < j_2 < \dots < j_n \leq L - 1$, then $\alpha(j_1)$ belongs to each conflict interval $X_{j_1}, X_{j_2}, \dots, X_{j_n}$, and the remaining part of the argument follows.

Therefore, $|I| = |I_{\alpha\beta}| + |I_{\beta\alpha}| \leq 256d^3(d + 2)^4 + 2$. ┘

We now consider an auxiliary undirected graph H where $V(H) = \{2, 3, \dots, L - 1\}$ and vertices i and j are connected by an edge if the conflict intervals X_i and X_j intersect. This is a well-known definition of an interval graph. Every interval graph is also a perfect graph [43], so H has a proper coloring where the number of colors is equal to the maximum size of a clique in H . In interval graphs, the maximum size of a clique is equal to the maximum number of intervals which intersect at a single point. This number is bounded by Claim 86 by $256d^3(d + 2)^4 + 2$, and therefore H can be colored using that many colors. Since $L = \lceil 256d^3(d + 2)^4 + 2 \rceil \ell + 2$, we have $|V(H)| = \lceil 256d^3(d + 2)^4 + 2 \rceil \ell$ and there exists a color that has been used for at least ℓ vertices of H . Therefore, H contains an independent set of size ℓ , or in other words, there exists a set of ℓ pairwise disjoint conflict intervals $X_{d_1}, X_{d_2}, \dots, X_{d_\ell}$ ($2 \leq d_1 < d_2 < \dots < d_\ell \leq L - 1$).

We create a modified identity ordered cage \mathcal{C}' which is a subset of \mathcal{C} given by indices d_1, d_2, \dots, d_ℓ .

Claim 87. \mathcal{C}' is a neighbor cage.

Proof. We fix an index $i \in \{2, 3, \dots, \ell - 1\}$, aiming to prove that b_{d_i} belongs to the region $\text{Area}(\mathcal{C}, d_{i-1}, d_{i+1})$. Naturally, $\text{Area}(\mathcal{C}, d_{i-1}, d_{i+1}) = \text{Area}(\mathcal{C}', i - 1, i + 1)$, so this is enough to complete the proof of the claim.

Lemma 82 asserts that $a_{d_i} \in \text{Int Area}(\mathcal{C}, d_i - 1, d_i + 1) \subseteq \text{Int Area}(\mathcal{C}, d_{i-1}, d_{i+1})$. Hence, $a_i \in \text{Int Area}(\mathcal{C}, 1, d_{i+1})$. Moreover, from Corollary 77 we deduce that $\text{Area}(\mathcal{C}, 1, d_{i-1}) \cap \text{Int Area}(\mathcal{C}, d_{i-1}, d_{i+1}) = \emptyset$. Hence, $a_i \notin \text{Area}(\mathcal{C}, 1, d_{i-1})$.

Now, we apply Claim 85. Since the conflict intervals $X_{d_1}, X_{d_2}, \dots, X_{d_\ell}$ are pairwise disjoint, we get that $\alpha(d_i) \notin X_{d_{i-1}}$ and $\alpha(d_i) \notin X_{d_{i+1}}$, and therefore a_{d_i} is separated from b_{d_i} by neither $\text{Area}(\mathcal{C}, 1, d_{i-1})$ nor $\text{Area}(\mathcal{C}, 1, d_{i+1})$. Due to the fact that $a_{d_i} \in \text{Int Area}(\mathcal{C}, d_{i-1}, d_{i+1})$, we get that $b_{d_i} \notin \text{Area}(\mathcal{C}, 1, d_{i-1})$ and $b_i \in \text{Area}(\mathcal{C}, 1, d_{i+1})$. Now, we observe that

$$\text{Area}(\mathcal{C}, 1, d_{i+1}) = \text{Area}(\mathcal{C}, 1, d_{i-1}) \cup \text{Area}(\mathcal{C}, d_{i-1}, d_{i+1})$$

(Lemma 76). Since $b_{d_i} \in \text{Area}(\mathcal{C}, 1, d_{i+1})$ and $b_{d_i} \notin \text{Area}(\mathcal{C}', 1, d_{i-1})$, we infer that $b_{d_i} \in \text{Area}(\mathcal{C}, d_{i-1}, d_{i+1})$. Hence, \mathcal{C}' is a neighbor cage of order ℓ . ┘

Claim 87 completes the proof. □

Let us now verify a few basic properties of neighbor cages.

Lemma 88. Fix a neighbor cage $\mathcal{C} = \text{NeighborCage}((a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, \mathcal{P}, \mathcal{Q})$ in a graph G embedded in the plane, and consider an edge $uv \in E(G)$. Assume that for some L, R ($2 \leq L < R < \ell$), we have that $u \in \text{Area}(\mathcal{C}, L, R)$, but $u \notin V(\mathcal{Q}) \cup \{p\}$. Then $v \in \text{Area}(\mathcal{C}, L - 1, R + 1)$.

Proof. Let $\mathcal{P} = \text{Tree}(P_1, P_2, \dots, P_\ell)$.

Since $\text{Area}(\mathcal{C}, L, R) \subseteq \text{Area}(\mathcal{C}, L - 1, R + 1)$ (a consequence of Lemma 76), we have that $u \in \text{Area}(\mathcal{C}, L - 1, R + 1)$. We now want to prove that $u \in \text{Int Area}(\mathcal{C}, L - 1, R + 1)$.

Assume on the contrary that $u \in \partial\text{Area}(\mathcal{C}, L - 1, R + 1)$. As $\text{Area}(\mathcal{C}, L, R) \subseteq \text{Area}(\mathcal{C}, L - 1, R + 1)$ and $u \in \text{Area}(\mathcal{C}, L, R)$, it follows that $u \in \partial\text{Area}(\mathcal{C}, L, R)$. Hence,

$$u \in [\partial\text{Area}(\mathcal{C}, L - 1, R + 1) \cap \partial\text{Area}(\mathcal{C}, L, R)] - [V(\mathcal{Q}) \cup \{p\}].$$

Since $\partial\text{Area}(\mathcal{C}, L, R) \subseteq V(\mathcal{P}) \cup V(\mathcal{Q})$, we infer from the definition of a neighbor area that

$$u \in \partial\text{Area}(\mathcal{C}, L, R) - V(\mathcal{Q}) \quad \Rightarrow \quad u \in \partial\text{Area}(\mathcal{C}, L, R) \cap V(\mathcal{P}) \quad \Rightarrow \quad u \in P_L \cup P_R.$$

Analogously, $u \in P_{L-1} \cup P_{R+1}$.

However, as \mathcal{P} is a simple geodesic tree, every pair of different paths P_i, P_j ($1 \leq i, j \leq \ell$) intersects only at p . Hence, $u = p$ — a contradiction, since we assumed that $u \neq p$. \square

Therefore, short paths can only reach very local parts of the neighbor cage without intersecting \mathcal{Q} or passing through p . This fact is explored in the following lemmas:

Lemma 89. *Fix a neighbor cage $\mathcal{C} = \text{NeighborCage}((a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, \mathcal{P}, \mathcal{Q})$ and two indices i, j ($1 \leq i < j \leq \ell$). No shortest path between b_i and a_j can pass through $\text{Root}(\mathcal{P})$.*

Proof. If any shortest path connecting b_i and a_j contains $\text{Root}(\mathcal{P})$ as a vertex, then

$$\begin{aligned} d < \text{dist}(b_i, a_i) &\leq \text{dist}(b_i, \text{Root}(\mathcal{P})) + \text{dist}(\text{Root}(\mathcal{P}), a_i) = \\ &= \text{dist}(b_i, \text{Root}(\mathcal{P})) + \text{dist}(\text{Root}(\mathcal{P}), a_j) = \text{dist}(b_i, a_j) \leq d \end{aligned}$$

— a contradiction. \square

Lemma 90. *Fix a neighbor cage $\mathcal{C} = \text{NeighborCage}((a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, \mathcal{P}, \mathcal{Q})$ in G and two indices i, j ($i < j$, $i \in [d + 2, \ell - d - 1]$). Take a shortest path R from b_i to a_j . Note that this path intersects \mathcal{Q} because $a_j \in V(\mathcal{Q})$. The first intersection of R with \mathcal{Q} belongs to $\text{Area}(\mathcal{C}, i - d - 1, i + d + 1)$.*

Proof. From the definition of a neighbor cage, we have that $b_i \in \text{Area}(\mathcal{C}, i - 1, i + 1)$. Let $R = (v_0, v_1, \dots, v_\delta)$ be the shortest path from b_i to a_j , where $v_0 = b_i$, $v_\delta = a_j$, and $\delta \leq d$ (as required by the definition of a semi-ladder). Let also v_k be the first vertex of this path belonging to $V(\mathcal{Q})$. Since $p \notin (v_0, v_1, \dots, v_\delta)$ (Lemma 89), by applying Lemma 88 we inductively deduce that $v_t \in \text{Area}(\mathcal{C}, i - t - 1, i + t + 1)$ for each $t \in \{0, 1, \dots, k\}$. Therefore, $v_k \in \text{Area}(\mathcal{C}, i - k - 1, i + k + 1)$. As $k \leq \delta \leq d$, we infer that $v_k \in \text{Area}(\mathcal{C}, i - d - 1, i + d + 1)$. \square

Now, for a vertex v belonging to a geodesic tree \mathcal{T} , we define the **depth** $\mu_{\mathcal{T}}(v)$ of v as the distance from v to $\text{Root}(\mathcal{T})$ in \mathcal{T} . Note that it is implied by the definition of a geodesic tree that $\mu_{\mathcal{T}}(v)$ is also equal to the distance between v and $\text{Root}(\mathcal{T})$ in G .

We remark a few simple facts about $\mu_{\mathcal{T}}$:

- $\mu_{\mathcal{T}}(\text{Root}(\mathcal{T})) = 0$.

- $\mu_{\mathcal{T}}$ is constant on the set of leaves of \mathcal{T} .
- If a vertex u is an ancestor of another vertex v in \mathcal{T} , then $\text{dist}(u, v) = \mu_{\mathcal{T}}(v) - \mu_{\mathcal{T}}(u)$.
- More generally, for any two vertices $u, v \in V(\mathcal{T})$, we have $\text{dist}(u, v) \geq |\mu_{\mathcal{T}}(u) - \mu_{\mathcal{T}}(v)|$.
- On any simple path in \mathcal{T} , the depths of vertices are first decreasing, and then increasing. In other words, there do not exist three vertices x, y, z lying on a simple path in \mathcal{T} in this order, for which $\mu_{\mathcal{T}}(x) < \mu_{\mathcal{T}}(y) > \mu_{\mathcal{T}}(z)$.

We already know from Lemmas 89 and 90 that each shortest path R connecting the vertices b_i and a_j of the semi-ladder ($i < j$) must avoid $\text{Root}(\mathcal{P})$ and the first intersection of R with \mathcal{Q} occurs inside of $\text{Area}(\mathcal{C}, i - d - 1, i + d + 1)$. It turns out that this first intersection imposes serious restrictions on each following intersection of R with \mathcal{Q} .

Lemma 91. *Fix a neighbor cage $\mathcal{C} = \text{NeighborCage}((a_i)_{i=1}^{\ell}, (b_i)_{i=1}^{\ell}, \mathcal{P}, \mathcal{Q})$ and two indices i, j ($1 \leq i < j \leq \ell$). Consider any shortest path from b_i to a_j : $R = (v_0, v_1, v_2, \dots, v_{\delta})$ for some $\delta \leq d$, $v_0 = b_i$, $v_{\delta} = a_j$. Suppose that for some indices $0 \leq x \leq y \leq \delta$ we have $v_x, v_y \in V(\mathcal{Q})$. Then $\mu_{\mathcal{Q}}(v_y) > \mu_{\mathcal{Q}}(\text{lca}_{\mathcal{Q}}(v_x, a_i))$.*

Proof. We set $\mu := \mu_{\mathcal{Q}}$, $\text{lca} := \text{lca}_{\mathcal{Q}}$.

For the sake of contradiction let us fix a shortest path R from b_i to a_j and two indices x, y as above such that $\mu(v_y) \leq \mu(\text{lca}(v_x, a_i))$. We consider a path G from b_i to a_i which first follows R from b_i to v_x , then goes up the tree \mathcal{Q} from v_x to $\text{lca}(v_x, a_i)$, and then goes down the tree \mathcal{Q} from $\text{lca}(v_x, a_i)$ to a_i (Figure 27).

We obviously have that

$$\text{dist}(b_i, a_i) \leq \text{len}(G) = \text{len}(G[b_i, v_x]) + \text{len}(G[v_x, \text{lca}(v_x, a_i)]) + \text{len}(G[\text{lca}(v_x, a_i), a_i]). \quad (8)$$

As G follows R on the fragment between b_i and v_x , we have that

$$\text{len}(G[b_i, v_x]) = \text{len}(R[b_i, v_x]). \quad (9)$$

Next, $\text{lca}(v_x, a_i)$ is an ancestor of v_x in \mathcal{Q} and thus

$$\begin{aligned} \text{len}(G[v_x, \text{lca}(v_x, a_i)]) &= \mu(v_x) - \mu(\text{lca}(v_x, a_i)) \leq \mu(v_x) - \mu(v_y) \leq \text{dist}(v_x, v_y) = \\ &= \text{len}(R[v_x, v_y]). \end{aligned} \quad (10)$$

Analogously, $\text{lca}(v_x, a_i)$ is an ancestor of a_i in \mathcal{Q} . Therefore,

$$\begin{aligned} \text{len}(G[\text{lca}(v_x, a_i), a_i]) &= \mu(a_i) - \mu(\text{lca}(v_x, a_i)) \leq \mu(a_i) - \mu(v_y) = \mu(a_j) - \mu(v_y) \leq \\ &\leq \text{dist}(v_y, a_j) = \text{len}(R[v_y, a_j]). \end{aligned} \quad (11)$$

By substituting (9), (10), (11) into (8), we conclude that

$$\begin{aligned} \text{dist}(b_i, a_i) &\leq \text{len}(R[b_i, v_x]) + \text{len}(R[v_x, v_y]) + \text{len}(R[v_y, a_j]) = \text{len}(R) = \text{dist}(b_i, a_j) = \\ &= \delta \leq d \end{aligned}$$

— a contradiction since the semi-ladder requires that $\text{dist}(b_i, a_i) > d$. □

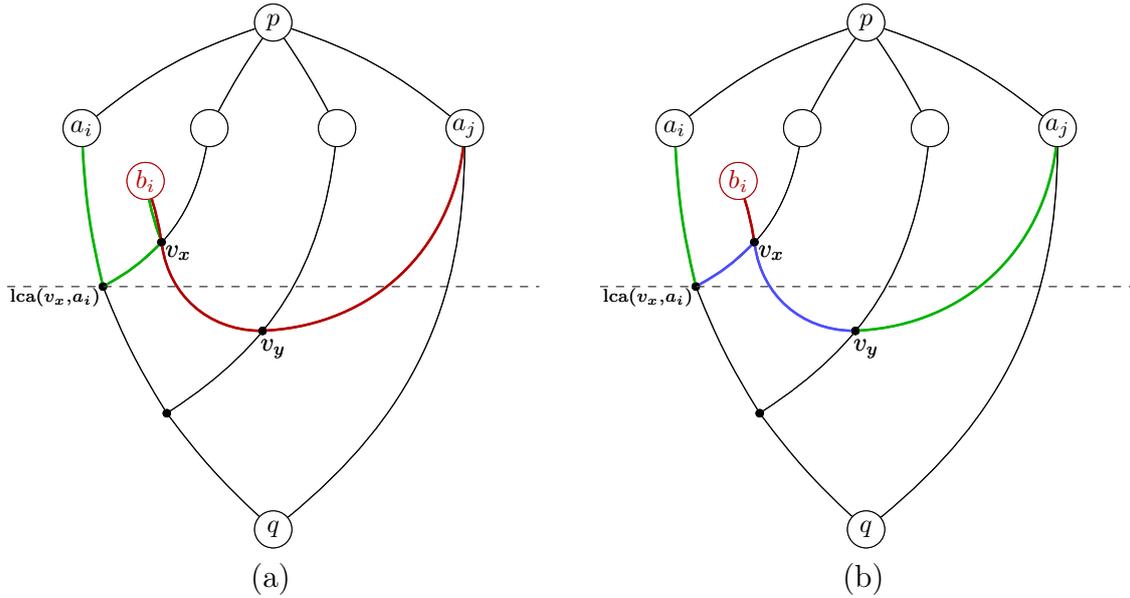


Figure 27: The setup in Lemma 91.

(a) — The shortest path R from b_i to a_j (red) crosses \mathcal{Q} twice. We are going to prove that there exists a path G from b_i to a_i (green) which is not longer than R .

(b) — Both R and G have been split into three parts: red, blue and green. R and G share the red fragment. (10) proves that the blue segment in G is not longer than the blue segment in R . (11) proves that the green segment in G is not longer than the green segment in R .

A less formal way of thinking about Lemma 91 is that as soon as the shortest path from b_i to a_j touches any vertex of \mathcal{Q} , all subsequent intersections of the path with \mathcal{Q} need to happen sufficiently deep in the tree. There are also some limitations on where this path may intersect \mathcal{Q} at all — by letting $x = y$ in the lemma above, we infer that each intersection v_x of a shortest path from b_i to a_j must satisfy $v_x \neq \text{lca}_{\mathcal{Q}}(v_x, a_i)$. That is, no intersection of this path with \mathcal{Q} can be an ancestor of a_i in \mathcal{Q} .

The most important conclusion will become clear in a moment: since the first intersection of every shortest path from b_i to a_j ($i < j$) with \mathcal{Q} must happen topologically relatively close to b_i (Lemma 90), we should be able to infer a common lower bound on the depth of every such intersection, no matter the path or value of j . This is indeed possible, albeit with slightly stronger assumptions, as proved in Lemma 93 below.

In order to proceed, we need a small definition.

Definition 92. In a neighbor cage $\mathcal{C} = \text{NeighborCage}((a_i)_{i=1}^{\ell}, (b_i)_{i=1}^{\ell}, \mathcal{P}, \mathcal{Q})$, a neighbor area $\text{Area}(\mathcal{C}, i, j)$ is **rootless** if $\text{Root}(\mathcal{Q}) \notin \text{Int Area}(\mathcal{C}, i, j)$.

Note that not all neighbor areas must be rootless, as seen in Figure 28.

We are now ready to introduce the promised lemma.

Lemma 93. Fix a neighbor cage $\mathcal{C} = \text{NeighborCage}((a_i)_{i=1}^{\ell}, (b_i)_{i=1}^{\ell}, \mathcal{P}, \mathcal{Q})$ and an index i ($i \in [d + 2, \ell - d - 1]$). If $\text{Area}(\mathcal{C}, i - d - 1, i + d + 1)$ is rootless, then each intersection v of a shortest path from b_i to a_j ($j > i$) with \mathcal{Q} satisfies

$$\mu_{\mathcal{Q}}(v) > \mu_{\mathcal{Q}}(\text{lca}_{\mathcal{Q}}(a_{i-d-1}, a_{i+d+1})).$$

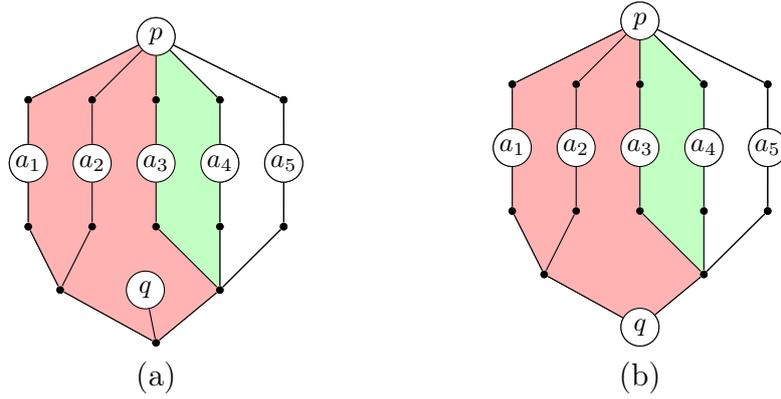


Figure 28: (a) — An example neighbor cage \mathcal{C} of order 5. $\text{Area}(\mathcal{C}, 3, 4)$ (green) is rootless, but $\text{Area}(\mathcal{C}, 1, 3)$ (red) is not. Vertices b_1, \dots, b_5 were omitted for clarity.
 (b) — Another neighbor cage \mathcal{D} . Now, both $\text{Area}(\mathcal{D}, 1, 3)$ and $\text{Area}(\mathcal{D}, 3, 4)$ are rootless.

Proof. We set $\mu := \mu_{\mathcal{Q}}$, $\text{lca} := \text{lca}_{\mathcal{Q}}$. We also let $v_{\text{lca}} := \text{lca}(a_{i-d-1}, a_{i+d+1})$ and $\mathcal{Q} = \text{Tree}(Q_1, \dots, Q_\ell)$.

Basing on the definition of neighbor areas, we infer that v_{lca} is the shallowest vertex in \mathcal{Q} belonging to $\partial \text{Area}(\mathcal{C}, i-d-1, i+d+1)$. We will now prove that v_{lca} is also the shallowest vertex in \mathcal{Q} belonging to $\text{Area}(\mathcal{C}, i-d-1, i+d+1)$. Assume for contradiction that there exists a vertex $x \in \text{Area}(\mathcal{C}, i-d-1, i+d+1) \cap V(\mathcal{Q})$ for which $\mu(x) < \mu(v_{\text{lca}})$. Since $\text{Root}(\mathcal{Q}) \notin \text{Int} \text{Area}(\mathcal{C}, i-d-1, i+d+1)$, the unique simple path in \mathcal{Q} connecting x with $\text{Root}(\mathcal{Q})$ must intersect $\partial \text{Area}(\mathcal{C}, i-d-1, i+d+1)$ at some vertex $s \in V(\mathcal{Q})$. Obviously, $\mu(s) \geq \mu(v_{\text{lca}})$. However, it means that the simple path in \mathcal{Q} connecting x with $\text{Root}(\mathcal{Q})$ contains three vertices in the following order: $x, s, \text{Root}(\mathcal{Q})$. But

$$\mu(x) < \mu(v_{\text{lca}}) \leq \mu(s) > \mu(\text{Root}(\mathcal{Q})) = 0 \quad \Rightarrow \quad \mu(x) < \mu(s) > \mu(\text{Root}(\mathcal{Q}))$$

— a contradiction. Hence, v_{lca} is the shallowest vertex of $\text{Area}(\mathcal{C}, i-d-1, i+d+1) \cap V(\mathcal{Q})$.

We fix a shortest path R from b_i to a_j . Let v_1 be the first intersection of R with $V(\mathcal{Q})$. By Lemma 90, we have that $v_1 \in \text{Area}(\mathcal{C}, i-d-1, i+d+1)$.

We now prove (Figure 29) that the simple path in \mathcal{Q} between v_1 and a_i — call it S — is fully contained within $\text{Area}(\mathcal{C}, i-d-1, i+d+1)$. Firstly, from the properties of neighbor cages, it follows that $a_i \in \text{Area}(\mathcal{C}, i-1, i+1)$; thus, $a_i \in \text{Area}(\mathcal{C}, i-d-1, i+d+1)$. Secondly, by the definition of neighbor areas, the intersection of $\partial \text{Area}(\mathcal{C}, i-d-1, i+d+1)$ with \mathcal{Q} is the simple path in \mathcal{Q} connecting $\text{Split}(\mathcal{C}, i-d-1)$ with $\text{Split}(\mathcal{C}, i+d+1)$. Let us call this path B . Therefore, if S was not contained within $\text{Area}(\mathcal{C}, i-d-1, i+d+1)$, it would have to intersect R at least twice — once when leaving the neighbor area, and once when reentering it. This would, however, mean that \mathcal{Q} contains a cycle — which is impossible.

Hence, S fully belongs to $\text{Area}(\mathcal{C}, i-d-1, i+d+1)$. It means in particular that

$$\begin{aligned} \mu(\text{lca}(v_1, a_i)) &= \min\{\mu(w) \mid w \in S\} \geq \\ &\geq \min\{\mu(w) \mid w \in \text{Area}(\mathcal{C}, i-d-1, i+d+1) \cap V(\mathcal{Q})\} = \mu(v_{\text{lca}}). \end{aligned}$$

Therefore, by Lemma 91, each intersection v of R with \mathcal{Q} — either v_1 , or any later

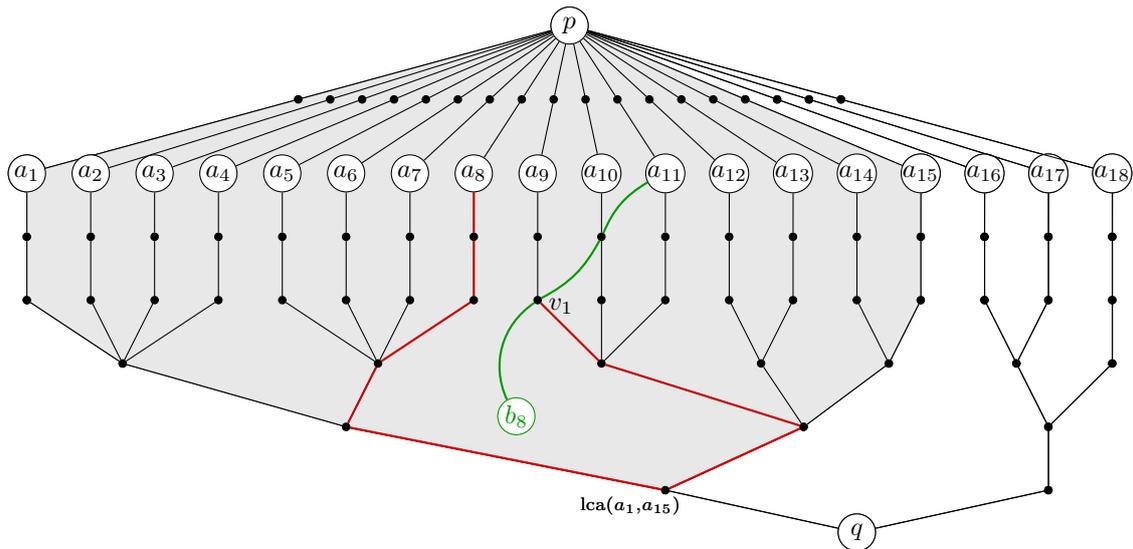


Figure 29: An example of the setup in Lemma 93 — a cage \mathcal{C} , $d = 6$ and $i = 8$. We consider v_1 — the first intersection of the shortest path between b_8 and a_{11} (green) with \mathcal{Q} . We prove that the simple path S in \mathcal{Q} connecting v_1 with a_8 (red) fully belongs to $\text{Area}(\mathcal{C}, 1, 15)$ (shaded). It will follow that $\mu(\text{lca}(v_1, a_8)) \geq \mu(v_{\text{lca}})$, and Lemma 91 will conclude the proof.

intersection — must satisfy

$$\mu(v) > \mu(\text{lca}(v_1, a_i)) \geq \mu(v_{\text{lca}}). \quad \square$$

5.5 Separating cages

Lemma 93 allows us to introduce the final variant of a cage, which will eventually lead us to a polynomial bound on the maximum order of any neighbor cage. This in turn will produce an upper bound on the maximum order of any distance- d semi-ladder.

Definition 94. A neighbor cage $\mathcal{C} = \text{NeighborCage}((a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, \mathcal{P}, \mathcal{Q})$ in a graph G is a **separating cage** if there exists an integer $\lambda \in \{0, 1, 2, \dots, d - 1\}$, called **threshold**, such that the following properties hold (Figure 30):

- for every two different indices i, j ($1 \leq i, j \leq \ell$), paths $\text{Path}(\mathcal{Q}, i)$ and $\text{Path}(\mathcal{Q}, j)$ do not share any vertices which are at depth larger than λ in \mathcal{Q} ;
- for every pair of indices i, j ($1 \leq i < j \leq \ell$), no shortest path from b_i to a_j contains any vertex belonging to \mathcal{Q} which is at depth at most λ in \mathcal{Q} .

We denote a separating cage in the following way:

$$\text{SeparatingCage}((a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, \mathcal{P}, \mathcal{Q}, \lambda).$$

We will now prove that graphs with large neighbor cages also contain large separating cages.

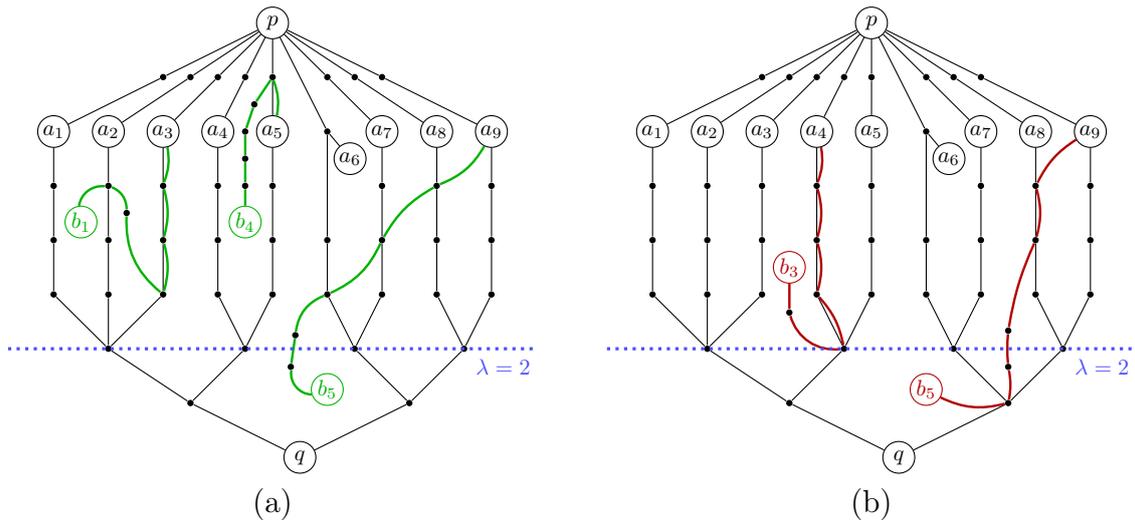


Figure 30: An example separating cage \mathcal{C} of order 9 and threshold $\lambda = 2$. The threshold is a blue dotted line; no two paths connecting $\text{Root}(\mathcal{Q})$ with a leaf of \mathcal{Q} can intersect above the line, and no path connecting b_i with a_j (for $i < j$) can intersect \mathcal{Q} at this line or below.

- (a) — Three valid paths connecting b_i with a_j for $i < j$ (green).
- (b) — Two invalid paths connecting b_i with a_j for $i < j$ (red). These cannot occur in the separating cage because each of them passes through a vertex of \mathcal{Q} at depth λ or less.

Lemma 95. *Each neighbor cage of order $2(2d + 3)((\ell - 1)d + 1)$ contains a separating cage of order ℓ as a subset.*

Proof. Let $L := 2(2d + 3)((\ell - 1)d + 1)$. Fix a cage $\mathcal{C} = \text{NeighborCage}((a_i)_{i=1}^L, (b_i)_{i=1}^L, \mathcal{P}, \mathcal{Q})$ of order L . We let $\mathcal{Q} = \text{Tree}(Q_1, Q_2, \dots, Q_L)$.

Let also $M := 2((\ell - 1)d + 1)$. We define a sequence (A_1, A_2, \dots, A_M) of M neighbor areas:

$$A_i = \text{Area}(\mathcal{C}, (2d + 3)(i - 1) + 1, (2d + 3)i) \quad \text{for each index } i \in [1, M].$$

We remark that because the set $\{(2d + 3)(i - 1) + 1, (2d + 3)i \mid i \in \{1, \dots, M\}\}$ contains only pairwise disjoint segments, we infer from Lemma 77 that the interiors of the neighbor areas A_1, A_2, \dots, A_M are pairwise disjoint. Hence, at least $M - 1 = 2(\ell - 1)d + 1$ out of neighbor areas A_1, A_2, \dots, A_M are rootless.

For each neighbor area A_i , we define its depth

$$\mu(A_i) = \mu_{\mathcal{Q}}(\text{lca}_{\mathcal{Q}}(a_{(2d+3)(i-1)+1}, a_{(2d+3)i})).$$

Note that $\mu(A_i)$ is chosen specifically to match the conditions of Lemma 93: provided that A_i is rootless, each intersection v of a shortest path from $b_{(2d+3)(i-1)+(d+2)}$ to a_j (for $j > (2d + 3)(i - 1) + (d + 2)$) with \mathcal{Q} satisfies $\mu_{\mathcal{Q}}(v) > \mu(A_i)$. We also remark that $\mu(A_i) \in \{0, 1, 2, \dots, d - 1\}$ for each $i \in [1, M]$.

The pigeonhole principle allows us to select $2\ell - 1$ out of $(2\ell - 2)d + 1$ rootless neighbor areas so that the selected neighbor areas have the same depth μ . We fix this subset of

rootless neighbor areas: $A_{t_1}, A_{t_2}, \dots, A_{t_{2\ell-1}}$ where $1 \leq t_1 < t_2 < \dots < t_{2\ell-1} \leq M$. Note that

$$\mu(A_{t_1}) = \mu(A_{t_2}) = \dots = \mu(A_{t_{2\ell-1}}) =: \lambda, \quad \text{Root}(\mathcal{Q}) \notin \text{Int Area}(A_{t_j}) \text{ for } j \in [1, 2\ell - 1].$$

We also define m_j for $j \in [1, 2\ell - 1]$ as the index of the central vertex of the semi-ladder in the neighbor area A_{t_j} : $m_j = (2d + 3)(t_j - 1) + (d + 2)$. Naturally, the sequence of indices m_j is increasing: $m_1 < m_2 < \dots < m_{2\ell-1}$. Also, for every pair of indices $i, j \in [1, 2\ell - 1]$ ($i < j$), each intersection v of a shortest path from b_{m_i} to a_{m_j} with \mathcal{Q} satisfies $\mu_{\mathcal{Q}}(v) > \lambda$.

Let \mathcal{C}' be the neighbor cage of order ℓ which is a subset of \mathcal{C} given by the sequence of indices $m_1, m_3, m_5, \dots, m_{2\ell-1}$:

$$\mathcal{C}' = \text{NeighborCage}((a_{m_1}, a_{m_3}, \dots, a_{m_{2\ell-1}}), (b_{m_1}, b_{m_3}, \dots, b_{m_{2\ell-1}}), \mathcal{P}', \mathcal{Q}'), \quad \mathcal{P}' \subseteq \mathcal{P}, \quad \mathcal{Q}' \subseteq \mathcal{Q}.$$

We verify that \mathcal{C}' is a separating cage with threshold λ .

Firstly, we see that $\lambda = \mu(B_{t_1}) \in \{0, 1, 2, \dots, d - 1\}$.

Secondly, we want to prove that no shortest path whose shortness is required by the semi-ladder in \mathcal{C}' contains any vertex belonging to \mathcal{Q} that is at depth at most λ in \mathcal{Q} . This is quickly resolved by Lemma 93. Fix any two vertices of the semi-ladder in \mathcal{C}' which require a short path: b_{m_i}, a_{m_j} ($i < j$, both i, j odd), and fix a shortest path R between these two vertices. As mentioned before, each intersection v of R with \mathcal{Q} satisfies $\mu_{\mathcal{Q}}(v) > \lambda$. Since $\mathcal{Q}' \subseteq \mathcal{Q}$, we obviously have $\mu_{\mathcal{Q}'}(v) > \lambda$ as well.

Finally, we want to prove that no two different paths connecting $\text{Root}(\mathcal{Q}')$ with the leaves of \mathcal{Q}' intersect at any vertex at depth larger than λ . We again fix two vertices of the semi-ladder in \mathcal{C}' : a_{m_i}, a_{m_j} ($i < j$, both i, j odd), and consider paths Q_{m_i} and Q_{m_j} , connecting $\text{Root}(\mathcal{Q})$ with a_{m_i} and a_{m_j} , respectively. Assume for contradiction that these paths intersect at a vertex v such that $\mu(v) > \lambda$. This means that $v \in Q_{m_i}$ and $v \in Q_{m_j}$. By the definition of identity ordered cages, we have that $v \in Q_k$ for every $k \in [m_i, m_j]$.

However, since both i and j are odd, we have that $i + 1 < j$. This means that

$$[m_{i+1} - d - 1, m_{i+1} + d + 1] \subseteq [m_i, m_j].$$

Hence, $v \in Q_{m_{i+1}-d-1}$ and $v \in Q_{m_{i+1}+d+1}$ since \mathcal{C} is an identity ordered cage. However, we verify that

$$\mu_{\mathcal{Q}}(\text{lca}_{\mathcal{Q}}(a_{m_{i+1}-d-1}, a_{m_{i+1}+d+1})) = \mu(A_{t_{i+1}}) = \lambda.$$

This means that paths $Q_{m_{i+1}-d-1}$ and $Q_{m_{i+1}+d+1}$ do not intersect at any vertex at depth greater than λ , yet v is an intersection of these paths at depth greater than λ — a contradiction. Therefore, paths Q_{m_i}, Q_{m_j} do not share any vertex at depth larger than λ . Since these paths were chosen arbitrarily from \mathcal{Q}' , all the paths connecting $\text{Root}(\mathcal{Q})$ with a leaf of \mathcal{Q}' are vertex-disjoint, as long as vertices at depth larger than λ are concerned.

Since all required properties of separating cages have been verified, \mathcal{C}' together with threshold λ is a separating cage. \square

In separating cages, we can prove a much more powerful variant of Lemma 88:

Lemma 96. *Consider a separating cage $\mathcal{C} = \text{SeparatingCage}((a_i)_{i=1}^{\ell}, (b_i)_{i=1}^{\ell}, \mathcal{P}, \mathcal{Q}, \lambda)$. For every pair of indices i, j ($1 \leq i < j \leq \ell$) and every edge uv on a shortest path between b_i and a_j , if $u \in \text{Area}(\mathcal{C}, L, R)$ for some integers $2 \leq L < R < \ell$, then $v \in \text{Area}(\mathcal{C}, L - 1, R + 1)$.*

Proof. In a similar vein to the proof of Lemma 88, we will prove that $u \notin \partial\text{Area}(\mathcal{C}, L - 1, R + 1)$; the correctness of the statement of the lemma will follow.

We remark that $u \in \text{Area}(\mathcal{C}, L, R)$, and thus $u \in \text{Area}(\mathcal{C}, L - 1, R + 1)$. We assume for contradiction that $u \in \partial\text{Area}(\mathcal{C}, L - 1, R + 1)$. Since $\text{Area}(\mathcal{C}, L, R) \subseteq \text{Area}(\mathcal{C}, L - 1, R + 1)$, we also have that $u \in \partial\text{Area}(\mathcal{C}, L, R)$. Therefore, both conditions below must be satisfied:

$$\begin{aligned} u \in \text{SPath}(\mathcal{C}, L - 1) \quad \vee \quad u \in \text{SPath}(\mathcal{C}, R + 1), \\ u \in \text{SPath}(\mathcal{C}, L) \quad \vee \quad u \in \text{SPath}(\mathcal{C}, R). \end{aligned}$$

Hence, v belongs to the intersection of $\text{SPath}(\mathcal{C}, i)$ and $\text{SPath}(\mathcal{C}, j)$ for $i \neq j$. Therefore, by Lemma 64, these paths intersect exactly at p and their common suffix belonging to \mathcal{Q} . However, since u lies on the shortest path whose shortness is required by the semi-ladder, Lemma 89 applies and asserts that $u \neq p$. It means that u belongs to the intersection of paths $\text{Path}(\mathcal{Q}, i)$ and $\text{Path}(\mathcal{Q}, j)$. But the definition of a separating cage requires that each such intersection must occur at depth at most λ . This means that the considered shortest path intersects \mathcal{Q} at a vertex u located at depth at most λ , which is explicitly forbidden by the definition of a separating cage — a contradiction.

As $u \in \text{Area}(\mathcal{C}, L - 1, R + 1)$, $u \notin \partial\text{Area}(\mathcal{C}, L - 1, R + 1)$, and v is connected to u by an edge, we have that $v \in \text{Area}(\mathcal{C}, L - 1, R + 1)$. \square

The lemma above leads to a straightforward linear bound on the maximum order of any separating cage:

Lemma 97. *Every separating cage has order smaller than $2d + 5$.*

Proof. Assume for contradiction that $\mathcal{C} = \text{SeparatingCage}((a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, \mathcal{P}, \mathcal{Q}, \lambda)$ is a separating cage of order $\ell \geq 2d + 5$. We consider a shortest path from b_{d+2} to a_{2d+4} : $(u_0, u_1, u_2, \dots, u_\delta)$ where $u_0 = b_{d+2}$, $u_\delta = a_{2d+4}$ and $\delta \leq d$ (as required by the semi-ladder).

By the properties of neighbor cages, we have that $u_0 \in \text{Area}(\mathcal{C}, d + 1, d + 3)$. Using Lemma 96, we prove inductively that $u_i \in \text{Area}(\mathcal{C}, d - i + 1, d + i + 3)$ for each $i \in \{0, 1, \dots, \delta\}$. In particular, we infer that $a_{2d+4} = u_\delta \in \text{Area}(\mathcal{C}, d - \delta + 1, d + \delta + 3)$. Since $\delta \leq d$, we get that $a_{2d+4} \in \text{Area}(\mathcal{C}, 1, 2d + 3)$.

However, Lemma 82 asserts that $a_{2d+4} \in \text{Int Area}(\mathcal{C}, 2d + 3, 2d + 5)$. Since

$$\text{Area}(\mathcal{C}, 1, 2d + 3) \cap \text{Int Area}(\mathcal{C}, 2d + 3, 2d + 5) = \emptyset$$

(Corollary 77), we have a contradiction. \square

We can now finalize the proof of Theorem 52 bounding the maximum distance- d semi-ladder order in planar graphs:

Corollary 98. *For $d \geq 1$, all distance- d semi-ladders in planar graphs have order smaller than*

$$d \cdot \left\{ d \cdot (2d - 1) \cdot \left(\left\{ [256d^3(d + 2)^4 + 2] \cdot 2(2d + 3)[(2d + 4)d + 1] + 1 \right\}^2 + 1 \right) + 2 \right\}^d + 1.$$

Proof. We define the following polynomial functions:

$$\begin{aligned}\chi_1(d) &= 2d + 5, \\ \chi_2(d) &= 2(2d + 3)[(\chi_1(d) - 1)d + 1], \\ \chi_3(d) &= [256d^3(d + 2)^4 + 2]\chi_2(d) + 2, \\ \chi_4(d) &= (\chi_3(d) - 1)^2 + 1, \\ \chi_5(d) &= (2d - 1)\chi_4(d).\end{aligned}$$

It can be verified that the formula in the statement of the corollary simplifies to

$$d \{d \cdot \chi_5(d) + 2\}^d + 1 =: M_d.$$

Assume that a distance- d semi-ladder of order at least M_d exists. We take its subset which is a semi-ladder of order M_d , and call it \mathcal{C}_1 .

Using Lemma 55, we find a quasi-cage \mathcal{C}_2 of order $\chi_5(d)$ in the graph.

Using Lemma 61, we find a cage \mathcal{C}_3 of order $\chi_4(d)$ as a subset of \mathcal{C}_3 .

Using Lemma 67, we associate the cage \mathcal{C}_3 with an order and create an ordered cage \mathcal{C}_4 underlying the same semi-ladder as \mathcal{C}_3 .

Using Lemma 73, we find an identity ordered cage \mathcal{C}_5 of order $\chi_3(d)$ as a subset of \mathcal{C}_4 . Note that this may require altering the embedding of the graph in the plane.

Using Lemma 84, we find a neighbor cage \mathcal{C}_6 of order $\chi_2(d)$ as a subset of \mathcal{C}_5 .

Using Lemma 95, we find a separating cage \mathcal{C}_7 of order $\chi_1(d) = 2d + 5$ as a subset of \mathcal{C}_6 .

Finally, Lemma 97 contradicts our assumption since it asserts that \mathcal{C}_7 cannot exist.

Hence, all distance- d semi-ladders in the class of planar graphs must have order smaller than M_d . \square

We remark that this upper bound is of the form $d \cdot \rho(d)^d + 1$ where $\rho(d)$ is a polynomial of degree 22, so the upper bound is of the form $d^{O(d)}$. This concludes the proof of Theorem 52.

6 Upper bounds in other classes of graphs

In this section we present upper bounds on the maximum semi-ladder orders in other classes of graphs: graphs with bounded maximum degree, graphs with bounded pathwidth or treewidth, and graphs excluding the complete graph K_t as a minor.

In Subsection 6.1, we prove the $\Delta^d + 1$ upper bound for the class of graphs with the maximum degree bounded by Δ . The proof is straightforward, since in these graphs, no vertex can have too many other vertices in its distance- d neighborhood.

In Subsection 6.2, we bound the maximum semi-ladder order from above in graphs with bounded pathwidth. As an important part of the proof, we generalize the sunflower lemma in a way that allows us to assign labels to the elements of each set of the sunflower.

In Subsection 6.3, we prove the $d^{O(d^{t-1})}$ upper bound for the class of graphs excluding K_t as a minor. This proof will utilize the elements of the standard toolchain of Sparsity — weak coloring numbers and uniform quasi-wideness — and can be generalized to any class of graphs with bounded expansion. However, in our proof, we also apply a new trick

specific to the considered class of graphs (Lemma 116), which allows us to prove a much tighter upper bound. The same section also concludes with a similar upper bound in the class of graphs with bounded treewidth, facilitated by Theorem 4.

6.1 Graphs of bounded degree

The following simple theorem presents an upper bound on the maximum semi-ladder order in graphs with maximum degree bounded by Δ . The main idea of the proof is that in a distance- d semi-ladder $a_1, a_2, \dots, a_\ell, b_1, b_2, \dots, b_\ell$, vertex b_1 must remain at distance at most d from vertices a_2, a_3, \dots, a_ℓ .

Theorem 99. *For each $\Delta \geq 2$, $d \geq 1$, the order of every distance- d semi-ladder in a graph with maximum degree bounded by Δ is bounded from above by $\Delta^d + 1$.*

Proof. Fix a graph G in which the maximum degree of a vertex is bounded from above by Δ , and let $a_1, a_2, \dots, a_\ell, b_1, b_2, \dots, b_\ell$ be a distance- d semi-ladder in G of order $\ell \geq 2$. By the properties of semi-ladders, we infer that for each $i \in \{2, 3, \dots, \ell\}$, we have that $\text{dist}(b_1, a_i) \leq d$ — that is, there are $\ell - 1$ vertices (different than b_1) in the distance- d neighborhood of b_1 . Since each vertex a_2, a_3, \dots, a_ℓ is different than b_1 (by the definition of a distance- d semi-ladder in Section 2), it means that each of them is at distance at least 1 and at most d from b_1 .

Obviously, there are at most $\Delta(\Delta - 1)^k$ vertices at distance exactly k ($k \geq 1$) from b_1 . Summing this over all $k \in \{1, 2, \dots, d\}$, we get that

$$\begin{aligned} \ell - 1 &\leq \sum_{k=1}^d \Delta(\Delta - 1)^{k-1} = \Delta + (\Delta - 1) \sum_{k=2}^d \Delta(\Delta - 1)^{k-2} \leq \Delta + (\Delta - 1) \sum_{k=1}^{d-1} \Delta^k = \\ &= 1 + (\Delta - 1) \sum_{k=0}^{d-1} \Delta^k = 1 + \frac{\Delta^d - 1}{\Delta - 1}(\Delta - 1) = \Delta^d. \quad \square \end{aligned}$$

We contrast the $\Delta^d + 1$ upper bound with the $\lfloor \frac{\Delta}{2} \rfloor^{\lceil \frac{d}{2} \rceil}$ lower bound (Theorem 5) proved in Section 3.

6.2 Graphs with bounded pathwidth

We begin this section by reminding about a classic combinatorial result, first proved by Erdős and Rado [11]: the **Sunflower Lemma**.

Definition 100. Suppose Ω is a universe and $\mathcal{F} = (F_1, F_2, \dots, F_n)$ is a family of (not necessarily different) subsets of Ω . Then \mathcal{F} is a **sunflower** if there exists a set C (which we call the **core** of the sunflower) such that for every two different indices $i, j \in [1, n]$, we have $F_i \cap F_j = C$. The **order** of a sunflower is the number of sets in it, that is, $|\mathcal{F}| = n$.

Theorem 101 (Sunflower Lemma). *Let \mathcal{G} be a family of $b!a^{b+1}$ (not necessarily different) subsets of Ω in which each subset has at most b elements. Then \mathcal{G} contains a subfamily $\mathcal{F} \subseteq \mathcal{G}$ which is a sunflower of order a .*

The “textbook proof” of this lemma can be found e.g. in the book by Cygan et al. [4].

In order to prove an upper bound on the maximum semi-ladder order in the class of graphs with bounded pathwidth, we first need to generalize the Sunflower Lemma.

Labeled Sunflower Lemma

We wish to generalize the Sunflower Lemma (Theorem 101) by allowing the sets in the family to assign **labels** to the elements of the set. We first need to describe formally what this labeling means.

Definition 102. For a finite set of labels Σ and a universe Ω , a **labeled subset** $A \subseteq \Omega$ over Σ is a subset of Ω in which each element is also assigned a label from Σ . The **cardinality** of A is the number of elements belonging to A .

In other words, a labeled subset is a partial function $A : \Omega \rightarrow \Sigma$, and the cardinality of A is the size of its domain.

The definition of a sunflower now naturally generalizes to the families of labeled subsets:

Definition 103. For a finite set of labels Σ and a universe Ω , a family \mathcal{F} of (not necessarily different) labeled subsets of Ω over Σ is a **labeled sunflower** if:

- \mathcal{F} is a sunflower (as in Definition 100), whose core we will call C , and
- for each element $v \in C$, each labeled subset $A \in \mathcal{F}$ assigns the same label to v ; that is, for each $v \in C$, there exists a label $\rho \in \Sigma$ such that $A(v) = \rho$ for each $A \in \mathcal{F}$.

It turns out that the textbook proof of the Sunflower Lemma [4] can be generalized to produce labeled sunflowers in a straightforward manner.

Lemma 104 (Labeled Sunflower Lemma). *For a finite set of labels Σ , a universe Ω and two integers $a \geq 1$, $b \geq 0$, any family of $ab!(a|\Sigma|)^b$ (not necessarily different) labeled subsets of Ω over Σ , each of cardinality at most b , contains a labeled sunflower of order a as a subset.*

Proof. We apply an induction on b . If $b = 0$, then the family contains at least $ab!(a|\Sigma|)^b = a$ empty sets and already is a labeled sunflower.

Now, assume that $b \geq 1$, and let \mathcal{F} be the family as in the statement of the lemma. We consider two cases:

- There exists an element $v \in \Omega$ belonging to at least $(b - 1)!(a|\Sigma|)^b$ labeled sets in \mathcal{F} . Hence, there exists a label $\rho \in \Sigma$ and a subfamily $\mathcal{F}' \subseteq \mathcal{F}$ such that we have $|\mathcal{F}'| \geq a(b - 1)!(a|\Sigma|)^{b-1}$, and each set $A \in \mathcal{F}'$ contains v and satisfies $A(v) = \rho$.

We create a modified universe $\Omega' := \Omega - \{v\}$, and form a family \mathcal{F}'' from \mathcal{F}' by removing v from each of the sets within \mathcal{F}' . Then, the Labeled Sunflower Lemma applies inductively to \mathcal{F}'' and parameters a and $b - 1$. Hence, we can find a labeled sunflower S of order a as a subset of \mathcal{F}'' . By reintroducing v to each of the sets in S , we find a labeled sunflower S' of order a as a subset of \mathcal{F}' .

- Each element $v \in \Omega$ occurs in fewer than $(b-1)!(a|\Sigma|)^b$ labeled sets in \mathcal{F} . Since $(b-1)!(a|\Sigma|)^b \leq \frac{1}{ab}|\mathcal{F}|$, we can find a family \mathcal{G} of a pairwise disjoint labeled sets in \mathcal{F} greedily. In each of a steps, we extend \mathcal{G} by any labeled set $A \in \mathcal{F}$, and remove from \mathcal{F} all labeled sets intersecting A . As each step removes at most $b!(a|\Sigma|)^b \leq \frac{1}{a}|\mathcal{F}|$ elements from \mathcal{F} , we can see that this is a valid construction of \mathcal{G} . Since the labeled sets in \mathcal{G} are pairwise disjoint, they form a labeled sunflower.

In both cases, we construct a labeled sunflower of order a as a subset of \mathcal{F} , so the proof is complete. \square

We remark that the lemma requires a relatively small multiplicative overhead on the size of the initial family of subsets, compared to the original statement of the lemma (Theorem 101):

$$\frac{ab!(a|\Sigma|)^b}{b!a^{b+1}} = |\Sigma|^b.$$

As a sidenote, there exists a simpler proof of a looser upper bound: find a sunflower of order $a|\Sigma|^b$ using the sunflower lemma, and then find a subset of this sunflower of order a which is a labeled sunflower (which is an easy task since there are at most $|\Sigma|^b$ labelings of the core of the sunflower). This proof, however, requires the multiplicative overhead equal to $|\Sigma|^{b(b+1)}$ compared to the Sunflower Lemma. If we used this version of the Labeled Sunflower Lemma later instead of the version we have just proved in Lemma 104, we would ultimately derive the $d^{O(p^2)}p^{O(p)}$ upper bound on the maximum distance- d semi-ladder size in graphs with pathwidth bounded by p , which is unfortunately too high for our needs.

Upper bound proof

In the first part of our proof, we will define a structure whose occurrence in a graph proves that the graph both contains a large semi-ladder and has low pathwidth.

Definition 105. In a graph G , for $p, d, k, \ell \geq 1$, a path decomposition (W_1, W_2, \dots, W_k) of width p , together with 2ℓ different vertices $a_1, a_2, \dots, a_\ell, b_1, b_2, \dots, b_\ell$, and ℓ different indices $t_1, t_2, \dots, t_\ell \in [1, k]$, is a **distance- d alignment** of order ℓ if:

- vertices $a_1, a_2, \dots, a_\ell, b_1, b_2, \dots, b_\ell$ form a distance- d semi-ladder in G ;
- vertices a_1, a_2, \dots, a_ℓ belong to the bags $W_{t_1}, W_{t_2}, \dots, W_{t_\ell}$, respectively.

We denote a distance- d alignment in the following way:

$$\text{Alignment}_d((W_i)_{i=1}^k, (a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, (t_i)_{i=1}^\ell).$$

We note that given a distance- d alignment $\text{Alignment}_d((W_i)_{i=1}^k, (a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, (t_i)_{i=1}^\ell)$, one can easily take a subset of the underlying semi-ladder, given by indices i_1, i_2, \dots, i_m such that $1 \leq i_1 < i_2 < \dots < i_m \leq \ell$, and form a distance- d alignment with order m : $\text{Alignment}_d((W_i)_{i=1}^k, (a_{i_j})_{j=1}^m, (b_{i_j})_{j=1}^m, (t_{i_j})_{j=1}^m)$.

In our proof, we will transform a graph with bounded pathwidth, and containing a large distance- d semi-ladder, into a distance- d alignment containing this semi-ladder. Then, we will find a reasonably large subset of this alignment having necessary structural properties. The following definition formalizes these properties:

Definition 106. For $p, d, k, \ell \geq 1$, a distance- d alignment of width p and order ℓ :

$$\text{Alignment}_d((W_i)_{i=1}^k, (a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, (t_i)_{i=1}^\ell)$$

is also a **distance- d sunflower alignment** if the following additional properties hold:

- bags $W_{t_1}, W_{t_2}, \dots, W_{t_\ell}$ form a sunflower, whose core we will denote by C ;
- in G , the distance- d profiles on C are identical for each vertex a_1, a_2, \dots, a_ℓ .

We will refer to distance- d sunflower alignment in the following way:

$$\text{SunflowerAlignment}_d((W_i)_{i=1}^k, (a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, (t_i)_{i=1}^\ell).$$

The set C is the **core** of the decomposition.

We are now ready to find large distance- d sunflower alignments in the graphs with large distance- d semi-ladders.

Lemma 107. For $p, d, \ell \geq 1$, every graph with pathwidth not exceeding p and containing a semi-ladder of order $\ell(p+1)![\ell(d+2)]^{p+1}$ also contains a distance- d sunflower path decomposition of order ℓ .

Proof. Let $L := \ell(p+1)![\ell(d+2)]^{p+1}$. Fix a graph G , together with its distance- d semi-ladder $a_1, a_2, \dots, a_L, b_1, b_2, \dots, b_L$. Since $\text{pw}(G) \leq p$, there exists a path decomposition $W = (W_1, W_2, \dots, W_k)$ of G of width at most p ; that is, we have $|W_i| \leq p+1$ for each $i \in [1, k]$. Moreover, we can assume that W is a **nice path decomposition**; that is, we have $W_1 = W_k = \emptyset$, and for each index $i \in \{2, 3, \dots, k\}$, either of the following conditions is satisfied:

- W_i **introduces** a vertex $v \in V(G)$; that is, $W_i = W_{i-1} \cup \{v\}$;
- W_i **forgets** a vertex $v \in V(G)$; that is, $W_i = W_{i-1} - \{v\}$.

It can be easily shown that each path decomposition can be converted into a nice path decomposition of the same width [4], so we can safely assume that W is a nice path decomposition.

For each vertex a_j ($1 \leq j \leq L$) of the semi-ladder, we can locate the bag W_{t_j} introducing a_j to the decomposition. Since the path decomposition is nice, the values t_j are different for all $j \in [1, L]$. Therefore, the following describes a distance- d alignment of order L :

$$\text{Alignment}_d((W_i)_{i=1}^k, (a_i)_{i=1}^L, (b_i)_{i=1}^L, (t_i)_{i=1}^L).$$

We will now find a subset of this distance- d alignment which is a distance- d sunflower alignment.

We first create a family \mathcal{F} of L labeled subsets of $V(G)$ over $\{0, 1, 2, \dots, d, +\infty\}$: for each $i \in [1, L]$, we let the i -th labeled subset in the family to be the bag W_{t_i} (containing the vertex a_i), where each vertex $v \in W_{t_i}$ has the label $\pi_d[a_i, W_{t_i}](v)$.

Secondly, since each set in \mathcal{F} contains at most $p+1$ elements, the cardinality of the set of labels is equal to $d+2$, and $|\mathcal{F}| = L = \ell(p+1)![\ell(d+2)]^{p+1}$, we infer from the

Labeled Sunflower Lemma (Lemma 104) that there exists a subfamily $\mathcal{F}' \subseteq \mathcal{F}$, which is a labeled sunflower of order ℓ .

Finally, we let $\mathcal{F}' = (W_{t_{i_1}}, W_{t_{i_2}}, \dots, W_{t_{i_\ell}})$ for $i_1 < i_2 < \dots < i_\ell$. We can now easily verify that the following structure is a distance- d sunflower alignment of order ℓ :

$$\text{SunflowerAlignment}_d \left((W_j)_{j=1}^k, (a_{i_j})_{j=1}^\ell, (b_{i_j})_{j=1}^\ell, (t_{i_j})_{j=1}^\ell \right). \quad \square$$

Now, it turns out that we can give an explicit upper bound on the order of every distance- d sunflower alignment:

Lemma 108. *For $d \geq 1$, no graph can contain a distance- d sunflower alignment of order $2d + 3$.*

Proof. Fix $\ell := 2d + 3$ and a graph G containing a distance- d sunflower alignment of order ℓ and width p :

$$\text{SunflowerAlignment}_d \left((W_i)_{i=1}^k, (a_i)_{i=1}^\ell, (b_i)_{i=1}^\ell, (t_i)_{i=1}^\ell \right).$$

We will prove that the shortest path connecting b_1 with a_i for some $i \in [2, \ell]$ has length greater than d , which contradicts the assumption that $a_1, a_2, \dots, a_\ell, b_1, b_2, \dots, b_\ell$ is a distance- d semi-ladder.

Let C be the core of the sunflower alignment. Firstly, we will prove that no shortest path from b_1 to a_i for $i \in [2, \ell]$ passes through C . Assume for contradiction that for some $i \in [2, \ell]$, the shortest path from b_1 to a_i contains a vertex $x \in C$ (Figure 31). Since $\text{dist}(x, a_i) \leq d$ and $\pi_d[a_1, C] = \pi_d[a_i, C]$, we infer that the vertex x is equidistant from a_1 and a_i . Hence,

$$\text{dist}(b_1, a_1) \leq \text{dist}(b_1, x) + \text{dist}(x, a_1) = \text{dist}(b_1, x) + \text{dist}(x, a_i) = \text{dist}(b_1, a_i) \leq d.$$

We have a contradiction as $\text{dist}(b_1, a_1) > d$ is required by the definition of a semi-ladder. We note that this fact means that, in particular, neither b_1 nor any of the vertices a_2, a_3, \dots, a_ℓ belongs to C .

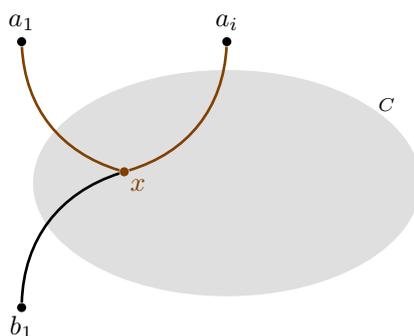


Figure 31: The setup in the proof that no shortest path from b_1 to a_i intersects C .

Next, we arrange the vertices a_2, a_3, \dots, a_ℓ in the order in which they appear in W . Formally, we define a permutation $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_{\ell-1})$ of $\{2, 3, \dots, \ell\}$ such that $t_{\sigma_1} < t_{\sigma_2} < \dots < t_{\sigma_{\ell-1}}$. This way, for each pair of indices i, j ($i, j \in [1, \ell - 1]$) with $i < j$,

the bag $W_{t_{\sigma_i}}$ containing a_{σ_i} appears earlier in W than the bag $W_{t_{\sigma_j}}$ containing a_{σ_j} (as $t_{\sigma_i} < t_{\sigma_j}$).

Let us find the vertex b_1 in the path decomposition. We already know that $b_1 \notin C$. Since $W_{t_{\sigma_{d+1}}} \cap W_{t_{\sigma_{d+2}}} = C$, we deduce that b_1 may belong to at most one of the bags $W_{t_{\sigma_{d+1}}}, W_{t_{\sigma_{d+2}}}$. As the definition of a path decomposition requires that b_1 belongs to the family of the bags forming a subinterval of the path decomposition, we get that at least one of the following cases is satisfied:

- (a) b_1 does not belong to any of the bags $W_1, W_2, W_3, \dots, W_{t_{\sigma_{d+1}}}$ of the decomposition;
- (b) b_1 does not belong any of the bags $W_{t_{\sigma_{d+2}}}, \dots, W_{k-2}, W_{k-1}, W_k$ of the decomposition.

As these cases are symmetrical, we assume without loss of generality that the case (a) holds. With this assumption, we consider the shortest path between b_1 and a_{σ_1} . We remark two facts about this path:

- it cannot pass through C , and
- it has to intersect each of the bags $W_{t_{\sigma_1}}, W_{t_{\sigma_2}}, \dots, W_{t_{\sigma_{d+1}}}$ of the decomposition (since $t_{\sigma_1} < t_{\sigma_2} < \dots < t_{\sigma_{d+1}}$, and $a_{\sigma_1} \in W_{t_{\sigma_1}}$).

Since each pair of the bags $W_{t_{\sigma_1}}, W_{t_{\sigma_2}}, \dots, W_{t_{\sigma_{d+1}}}$ intersects exactly at C , and the path cannot pass through C , we conclude that the path has to intersect each $W_{t_{\sigma_1}}, W_{t_{\sigma_2}}, \dots, W_{t_{\sigma_{d+1}}}$ at a different vertex. By our assumption that case (a) holds, b_1 does not belong to any of these bags. We conclude that the path must contain at least $d + 2$ vertices, hence its length is at least $d + 1$ — a contradiction.

This contradiction completes the proof of the lemma. □

This fact immediately allows us to conclude with a concrete upper bound.

Theorem 109. *For $p \geq 1$, $d \geq 1$, no graph with pathwidth bounded from above by p contains a distance- d semi-ladder of order*

$$(2d + 3)(p + 1)![(2d + 3)(d + 2)]^{p+1}.$$

Proof. If a graph with pathwidth bounded by p contained a distance- d semi-ladder of this order, it would also contain a distance- d sunflower alignment of order $2d + 3$ (Lemma 107). However, by Lemma 108, such an alignment cannot exist. □

We also remark the asymptotic version of this bound.

Corollary 110. *There exists a constant $\delta > 0$ such that for each $p, d \geq 2$, the maximum order of a distance- d semi-ladder in the class of graphs with pathwidth bounded by p is bounded from above by $(pd)^\delta$.*

We recall the construction attesting the $d^{\Omega(p)}$ lower bound in graphs with pathwidth bounded by p (Corollary 11).

6.3 Graphs with bounded treewidth and minor-free graphs

In this part, we will prove an upper bound on the maximum order of a distance- d semi-ladder in the class of graphs excluding K_t as a minor; an analogous bound for graphs with bounded treewidth will follow.

We begin by introducing two concepts of Sparsity that will become useful in the proof: **weak coloring numbers** and **uniform quasi-wideness**.

Definition 111 (weak coloring numbers). We follow the definition by Zhu [44].

Let $\Sigma_{V(G)}$ denote the set of all permutations of $V(G)$. For a permutation $\sigma \in \Sigma_{V(G)}$ and two vertices $u, v \in V(G)$, we write $u <_\sigma v$ if u occurs in σ before v .

For a fixed distance $d \geq 1$, we say that v is **weakly d -reachable** from u with respect to $\sigma \in \Sigma_{V(G)}$ if $v \leq_\sigma u$ and there exists a path of length at most d connecting u and v such that v is the minimum vertex on this path with respect to σ . We remark that u is weakly d -reachable from itself for every $d \in \mathbb{N}$. We denote the set of vertices weakly d -reachable from u with respect to σ as $\text{WReach}_d[G, \sigma, u]$ (Figure 32).

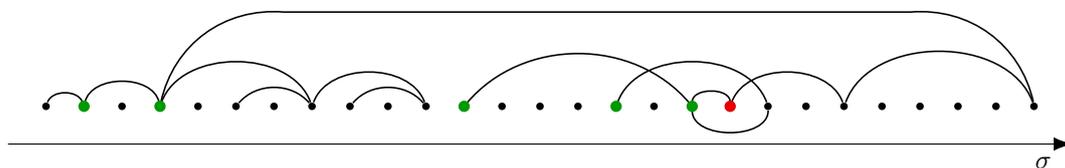


Figure 32: An example setup. In a permutation σ , for a vertex u (red), $\text{WReach}_4[G, \sigma, u] - \{u\}$ is marked green.

Given a graph G and an order $\sigma \in \Sigma_{V(G)}$ we define the weak coloring number of G with respect to σ as

$$\text{wcol}_d(G, \sigma) := \max_{u \in V(G)} |\text{WReach}_d[G, \sigma, u]|.$$

Finally, given a graph G , we define the **weak coloring number** of G as

$$\text{wcol}_d(G) := \min_{\sigma \in \Sigma_{V(G)}} \text{wcol}_d(G, \sigma).$$

Definition 112 (uniform quasi-wideness). A class \mathcal{C} of graphs is **uniformly quasi-wide** with margins $N : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ and $s : \mathbb{N} \rightarrow \mathbb{N}$ if for every $m, d \geq 1$, graph $G \in \mathcal{C}$ and a subset A of vertices such that $|A| \geq N(m, d)$, there exists a set of vertices $S \subseteq V(G)$, $|S| \leq s(d)$ and a subset $B \subseteq A - S$, $|B| \geq m$, which is a distance- d independent set in $G - S$.

It turns out that a subgraph-closed class \mathcal{C} of graphs is uniformly quasi-wide for some margins if and only if it is nowhere dense [27]. For more properties and applications of uniform quasi-wideness, we refer to the lecture notes from the University of Warsaw [33].

In this section, we may provide the margins N, s as partial functions, defined only for distances d from an infinite subset $D \subsetneq \mathbb{N}$. It is straightforward to see that these margins can be easily extended to total functions: given $m, d \in \mathbb{N}$, we set $d' := \min\{x \in D \mid x \geq d\}$, and we put $N(m, d) := N(m, d')$ and $s(d) := s(d')$.

In the proof of the upper bound for K_t -minor-free graphs, we will use the upper bound by Fabiański et al. [12], applicable to every uniformly quasi-wide class of graphs:

Lemma 113 ([12, Lemma 29]). *In every uniformly quasi-wide class of graphs \mathcal{C} with margins $N : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ and $s : \mathbb{N} \rightarrow \mathbb{N}$, each semi-ladder in every graph $G \in \mathcal{C}$ has its order bounded by $N(2 \cdot (d + 2)^{s(2d)} + 1, 2d)$.*

We will now prove that the class of K_t -minor-free graphs is uniformly quasi-wide by utilizing the generalized coloring numbers. Precisely, we will use the following upper bound on the weak coloring numbers in this class, proved by van den Heuvel et al. [42].

Lemma 114 ([42]). *For $t \geq 4$ and any graph G which excludes K_t as a minor, we have that $\text{wcol}_d(G) \leq \binom{d+t-2}{t-2}(t-3)(2d+1)$.*

We only need to prove that if a class has fairly low weak coloring numbers, then it is also uniformly quasi-wide with reasonably low margins. We remark that a very similar bound was shown in a work of Nadara et al. [21, Theorem 4.4], however the proof presented here is quite different and may be of independent interest.

Lemma 115. *Assume that a class of graphs \mathcal{C} has a finite weak coloring number $\text{wcol}_d(\mathcal{C})$ for each $d \geq 1$. Then \mathcal{C} is also uniformly quasi-wide with margins*

$$N(m, d) = (\text{wcol}_d(\mathcal{C}))! (m + 1)^{\text{wcol}_d(\mathcal{C})+1}, \quad s(d) = \text{wcol}_d(\mathcal{C}) \quad \text{for each } d \geq 1, m \geq 2.$$

Proof. Fix integers $d \geq 1, m \geq 2$, a graph $G \in \mathcal{C}$, and a subset A of the vertices of G containing at least $N(m, d)$ vertices. We will prove the lemma by showing that we can erase at most $\text{wcol}_d(\mathcal{C})$ vertices from G in order to find a distance- d independent set $B \subseteq A$ of size m in the remaining part of the graph.

In G , we can find an ordering σ of its vertices for which $\text{wcol}_d(G, \sigma) \leq \text{wcol}_d(\mathcal{C})$. Knowing that, we define a family \mathcal{F} of subsets of vertices:

$$\mathcal{F} := \{R_v \mid v \in S\} \quad \text{where } R_v := \text{WReach}_d[G, \sigma, v] \text{ for } v \in S.$$

Since each set in the family contains at most $\text{wcol}_d(\mathcal{C})$ elements and $|\mathcal{F}| \geq N(m, d) = (\text{wcol}_d(\mathcal{C}))! (m + 1)^{\text{wcol}_d(\mathcal{C})+1}$, the Sunflower Lemma (Theorem 101) lets us find a subfamily $\mathcal{F}' \subseteq \mathcal{F}$ which is a sunflower of order $m + 1$. We let $\mathcal{F}' = \{R_{v_1}, R_{v_2}, \dots, R_{v_{m+1}}\}$ for $m + 1$ different vertices v_1, v_2, \dots, v_{m+1} . We denote the core of \mathcal{F}' as D .

Without loss of generality, we assume that the vertices v_1, v_2, \dots, v_{m+1} are ordered decreasingly with respect to the order σ : $v_{m+1} <_\sigma v_m <_\sigma \dots <_\sigma v_1$. By the definition of weak reachability in Definition 111, for each $i \in [1, m + 1]$, v_i is the maximum element of R_{v_i} with respect to σ . Since for $i \in [1, m]$ we have $v_i >_\sigma v_{m+1}$, and the core D must be a subset of $R_{v_{m+1}}$, we infer that $v_i \notin D$ for each $i \in [1, m]$. Therefore, the set $\{v_1, v_2, \dots, v_m\}$ of vertices is disjoint with D . We now focus on the sunflower $\mathcal{F}'' := \{R_{v_1}, R_{v_2}, \dots, R_{v_m}\}$ of order m , whose core is naturally D . Obviously, $|D| \leq \text{wcol}_d(\mathcal{C})$.

We define the graph $G' := G - D$. We will now prove that $\{v_1, v_2, \dots, v_m\}$ is a distance- d independent set in G' . Assume for contradiction that there exists a path of length at most d connecting two different vertices of the set — say, v_i, v_j for two different indices $i, j \in [1, m]$. Let w_{\min} be the vertex on this path which is minimal with respect to σ . This, however, implies that w_{\min} belongs to both $\text{WReach}_d[G, \sigma, v_i]$ and $\text{WReach}_d[G, \sigma, v_j]$, and thus $w_{\min} \in R_{v_i} \cap R_{v_j}$. Since $\mathcal{F}'' = \{R_{v_1}, R_{v_2}, \dots, R_{v_m}\}$, we infer that w_{\min} belongs to D — the core of \mathcal{F}'' . This is a contradiction since D is disjoint with G' . \square

We remark that a combination of Lemmas 113, 114 and 115 can already provide us a concrete upper bound on the maximum order of a distance- d semi-ladder in the class of K_t -minor-free graphs. However, the produced bound will be insufficient for our purposes — it can be verified that it would result in the $d^{O(d^{2t-2})}$ upper bound for each fixed $t \geq 4$, while we aim at the $d^{O(d^{t-1})}$ bound.

In order to prove a tighter upper bound, we will improve the margins N, s with which the class of K_t -minor-free graphs is uniformly quasi-wide using the lemma below. This lemma is inspired by Lemma 3.16 from the work of Nešetřil and Ossona de Mendez [27]. The ideas from the original lemma have been adapted in order to show a bound with the margin s dependent only on t (and not on the radius d), and with the margin N still relatively low. We note that an idea of proving uniform quasi-wideness with the margin s independent on d in minor-free classes of graphs has been presented independently in a work of Pilipczuk et al. [35].

Lemma 116. *Fix $t \geq 4$. Assume that a class of graphs \mathcal{C} :*

- *excludes the complete bipartite graph $K_{t-1, t-1}$ as a minor, and*
- *is uniformly quasi-wide with margins $N : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$, $s : \mathbb{N} \rightarrow \mathbb{N}$.*

Then \mathcal{C} is also uniformly quasi-wide with margins $\widehat{N} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ and $\widehat{s} : \mathbb{N} \rightarrow \mathbb{N}$, where \widehat{s} is a constant function, $\widehat{s} \equiv t - 2$, and

$$\widehat{N}(m, 2d) = N \left(\binom{s(2d) + 1}{t - 1} (m + t), 2d \right) \quad \text{for each } m, d \geq 1.$$

Proof. We fix a graph $G \in \mathcal{C}$, integers $m, d \geq 1$, and a set of vertices $A \subseteq V(G)$ of size at least $\widehat{N}(m, 2d)$. By our assumption, there exists a subset $S \subseteq V(G)$ with $|S| \leq s(2d)$ and a subset $B \subseteq A - S$ with $|B| \geq \binom{s(2d)+1}{t-1} (m+t)$ such that B is a distance- $2d$ independent set in $G - S$.

Similarly to [27, Lemma 3.16], for each $v \in B$, we let L_v to be a minimal subset of S such that in $G - L_v$, there is no path of length d or less from v to any vertex from $S - L_v$; such a subset exists since S itself satisfies this requirement. Now, we will define a new set L'_v in the following way: if $|L_v| \leq t - 2$, then we set $L'_v := L_v$; otherwise, we take L'_v to be any subset of L_v of size exactly $t - 1$. Since $L'_v \subseteq S$, $|L'_v| \leq t - 1$ and $|S| \leq s(2d)$, we conclude that there are at most

$$\binom{s(2d)}{0} + \binom{s(2d)}{1} + \dots + \binom{s(2d)}{t-1} = \binom{s(2d) + 1}{t-1}$$

different possible sets L'_v . It follows that for some set $C \subseteq S$, $|C| \leq t - 1$, there are at least $m + t$ different vertices v_1, v_2, \dots, v_{m+t} for which $L'_{v_1} = L'_{v_2} = \dots = L'_{v_{m+t}} = C$.

We consider two cases, distinguished by the cardinality of C :

- If $|C| \leq t - 2$, then $L_{v_1} = L_{v_2} = \dots = L_{v_{m+t}} = C$. Now, assume there are two different vertices v_i, v_j ($i, j \in [1, m + t]$) at distance at most $2d$ in $G - C$. Since v_i and v_j are at distance greater than $2d$ in $G - S$, the shortest path between these two vertices in $G - C$ must pass through some vertex $x \in S - C$. This, however,

means that x is at distance at most d from one of the vertices v_i, v_j in $G - C$ — a contradiction since $x \notin L_{v_i}$ and $x \notin L_{v_j}$.

Therefore, $\{v_1, v_2, \dots, v_{m+t}\}$ is a distance- $2d$ independent set in $G - C$, where $|C| \leq t - 2$.

- If $|C| = t - 1$, then we will find $K_{m+t, t-1}$ as a minor in G , which will lead to a contradiction with $G \in \mathcal{C}$. Namely, for each $x \in \{v_1, v_2, \dots, v_{m+t}\}$, we construct a tree Y_x rooted at x with depth not exceeding d whose set of leaves is L_x and that does not contain any other vertices of S . Such a tree exists as the minimality of L_x implies that for each $v \in L_x$, the distance between x and v in $G - (L_x - \{v\})$ does not exceed d .

We see that no pair of trees Y_x, Y_y can intersect at any vertex outside of S — otherwise, we would be able to construct a path of length at most $2d$ between x and y avoiding S . This is, of course, impossible as $x, y \in A$.

Finally, we can find $K_{m+t, t-1}$ as a minor of G by contracting all subgraphs $Y_x - S$ into single vertices, and then observing that for each $x \in \{v_1, v_2, \dots, v_{m+t}\}$ and $y \in C$, the vertex y is adjacent to the contracted subgraph $Y_x - S$. We have a contradiction since we assumed G to be $K_{t-1, t-1}$ -free. \square

We remark that Lemma 116 also applies if \mathcal{C} excludes K_t as a minor for some fixed $t \geq 4$, for such classes of graphs exclude $K_{t-1, t-1}$ as a minor as well.

With our additional lemma in hand, we can prove the following upper bound:

Theorem 117. *There exists a polynomial p such that for $t \geq 4$, $d \geq 2$, there is no distance- d semi-ladder of order larger than*

$$d^{p(t) \cdot (2d+1)^{t-1}}$$

in any K_t -minor-free graph.

Proof. In this proof, $\text{poly}(t)$ will denote any polynomial of t with constant coefficients independent on d or t . However, each occurrence of this term in any equation might mean a different polynomial.

Fix $t \geq 4$ and $d \geq 2$, and consider the class \mathcal{C} of graphs excluding K_t as a minor. We first bound the weak coloring number of \mathcal{C} using Theorem 114:

$$\begin{aligned} \text{wcol}_d(\mathcal{C}) &\leq \binom{d+t-2}{t-2} (t-3)(2d+1) = \frac{(d+1)(d+2) \cdots (d+t-2)}{(t-2)!} (t-3)(2d+1) = \\ &= \frac{d+1}{1} \cdot \frac{d+2}{2} \cdots \frac{d+t-2}{t-2} \cdot (t-3)(2d+1) \leq \\ &\leq (d+1)^{t-2} (t-3)(2d+1) \leq \text{poly}(t) \cdot (d+1)^{t-1}. \end{aligned} \tag{12}$$

By Lemma 115, \mathcal{C} is uniformly quasi-wide with margins

$$\begin{aligned} N(m, d) &= (\text{wcol}_d(\mathcal{C}))! (m+1)^{\text{wcol}_d(\mathcal{C})+1} \leq [(m+1) \text{wcol}_d(\mathcal{C})]^{\text{wcol}_d(\mathcal{C})+1} = \\ &\stackrel{(12)}{=} [m \cdot \text{poly}(t) \cdot (d+1)^{t-1}]^{\text{poly}(t) \cdot (d+1)^{t-1}} \leq (md)^{\text{poly}(t) \cdot (d+1)^{t-1}} \end{aligned} \tag{13}$$

and

$$s(d) = \text{wcol}_d(\mathcal{C}) \stackrel{(12)}{\leq} \text{poly}(t) \cdot (d+1)^{t-1}. \quad (14)$$

In order to apply Lemma 116, we first bound the following expression for $m \geq 1$:

$$\begin{aligned} \binom{s(2d)+1}{t-1} (m+t) &\leq [s(2d)+1]^{t-1} (m+t) \stackrel{(14)}{\leq} [\text{poly}(t) \cdot (2d+1)]^{(t-1)^2} (m+t) \leq \\ &\leq m \cdot [d \cdot \text{poly}(t)]^{(t-1)^2}. \end{aligned} \quad (15)$$

Now, we get that \mathcal{C} is also uniformly quasi-wide with margins $\widehat{s} \equiv t-2$ and

$$\begin{aligned} \widehat{N}(m, 2d) &= N\left(\binom{s(2d)+1}{t-1} (m+t), 2d\right) \stackrel{(15)}{\leq} N\left(m \cdot [d \cdot \text{poly}(t)]^{(t-1)^2}, 2d\right) \leq \\ &\stackrel{(13)}{\leq} \left\{m \cdot [d \cdot \text{poly}(t)]^{(t-1)^2} \cdot 2d\right\}^{\text{poly}(t) \cdot (2d+1)^{t-1}} \leq (md)^{\text{poly}(t) \cdot (2d+1)^{t-1}}. \end{aligned} \quad (16)$$

Finally, we apply Lemma 113 and infer that the maximum order of any distance- d semi-ladder is bounded from above by

$$\widehat{N}\left(2 \cdot (d+2)^{t-2} + 1, 2d\right) \stackrel{(16)}{\leq} d^{\text{poly}(t) \cdot (2d+1)^{t-1}}.$$

Hence, there exists a polynomial p for which this maximum order is bounded from above by $d^{p(t) \cdot (2d+1)^{t-1}}$ for all $t \geq 4$, $d \geq 2$. \square

We also infer the following asymptotic version of Theorem 117:

Corollary 118. *For each fixed $t \geq 4$, the maximum order of a distance- d semi-ladder in the class of graphs excluding K_t as a minor is bounded from above by $d^{O(d^{t-1})}$.*

We contrast this result with the $2^{d^{\Omega(t)}}$ lower bound (Corollary 25) in the class of K_t -minor-free graphs.

By Theorem 4 (each graph G with $\text{tw}(G) \leq t$ is also K_{t+2} -minor-free), we immediately prove the following result from Corollary 118:

Corollary 119. *For each fixed $t \geq 2$, the maximum order of a semi-ladder in the class of graphs with treewidth not exceeding t is bounded from above by $d^{O(dt+1)}$.*

Similarly, we recall the $2^{d^{\Omega(t)}}$ lower bound on the maximum order of distance- d semi-ladders in the class of graphs with treewidth bounded by t (Corollary 23).

7 Conclusions

In this work we gave asymptotic lower and upper bounds on the maximum orders of distance- d half graphs and semi-ladders that can appear in various classes of sparse graphs: graphs with bounded maximum degree, planar graphs, graphs with bounded pathwidth or treewidth, and graphs excluding K_t as a minor. All these bounds are asymptotically almost tight and allow us to determine how large structures of this type may appear in graphs from these classes. Graphs with degree bounded by Δ and planar

graphs admit exponentially large distance- d half graphs. Graphs with pathwidth bounded by p admit polynomially large half graphs, with the degree of the polynomial depending linearly on p . Finally, the maximum order of distance- d half graphs in graphs with treewidth bounded by t and in K_t -minor-free graphs grows exponentially on d , where the exponent is a polynomial of degree depending linearly on t .

There are a few natural open problems regarding half graphs and semi-ladders in sparse classes of graphs and related topics:

- In case of the maximum order of distance- d semi-ladders in planar graphs, it would be interesting to close the asymptotic gap between $2^{\Omega(d)}$ and $d^{O(d)}$.
- We can easily see that in an arbitrary class \mathcal{C} with bounded expansion, the combination of Lemmas 113 and 115 results in a $(d \cdot \text{wcol}_{2d}(\mathcal{C}))^{O(\text{wcol}_{2d}(\mathcal{C})^2)}$ upper bound on the maximum order of any distance- d semi-ladder. However, we can deduce from Subsection 6.3 that as soon as \mathcal{C} excludes a fixed graph as a minor, we can derive a much tighter $(d \cdot \text{wcol}_{2d}(\mathcal{C}))^{O(\text{wcol}_{2d}(\mathcal{C}))}$ upper bound thanks to Lemma 116. Does the refined bound apply for all classes \mathcal{C} with bounded expansion?
- Can we characterize the classes \mathcal{C} of graphs where the orders of distance- d semi-ladders are bounded polynomially in d ? For instance, any class of graphs with bounded pathwidth has this property, but this is not the case for planar graphs or even for graphs with maximum degree bounded by 4.
- Theorem 26 proves that the neighborhood complexity in planar graphs is bounded by a polynomial of the size of the set $|X|$ and the distance d . The proof of Reidl et al. [38] provides a bound which is linear in $|X|$, but exponential in d . Can these bounds be reconciled? That is, can we bound the neighborhood complexity in planar graphs by a function which is linear in $|X|$ and polynomial in d ?

Acknowledgements

We would like to thank Michał Pilipczuk for pointing us to this problem, the verification of the correctness of the proofs in this work, and proofreading this work. Also, we want to thank Piotr Micek for reviewing this work and valuable suggestions regarding Section 4.

Last but not least, we wish to thank an anonymous reviewer for their helpful comments on the manuscript.

References

- [1] H. Adler and I. Adler. Interpreting nowhere dense graph classes as a classical notion of model theory. *Eur. J. Comb.*, 36:322–330, 2014.
- [2] G. Agnarsson and R. Greenlaw. *Graph Theory: Modeling, Applications, and Algorithms*. Pearson/Prentice Hall, 2007.
- [3] B. S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *J. ACM*, 41(1):153–180, 1994.
- [4] M. Cygan, F. V. Fomin, Ł. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, 2015.

- [5] A. Dawar and S. Kreutzer. Domination problems in nowhere-dense classes. In R. Kannan and K. N. Kumar, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2009, December 15-17, 2009, IIT Kanpur, India*, volume 4 of *LIPICs*, pages 157–168. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2009.
- [6] M. Debski, S. Felsner, P. Micek, and F. Schröder. Improved bounds for centered colorings. In S. Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2212–2226. SIAM, 2020.
- [7] P. G. Drange, M. S. Dregi, F. V. Fomin, S. Kreutzer, D. Lokshtanov, M. Pilipczuk, M. Pilipczuk, F. Reidl, F. S. Villaamil, S. Saurabh, S. Siebertz, and S. Sikdar. Kernelization and sparseness: the case of dominating set. In N. Ollinger and H. Vollmer, editors, *33rd Symposium on Theoretical Aspects of Computer Science, STACS 2016, February 17-20, 2016, Orléans, France*, volume 47 of *LIPICs*, pages 31:1–31:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- [8] L. Dubois, G. Joret, G. Perarnau, M. Pilipczuk, and F. Pitois. Two lower bounds for p -centered colorings. *Discret. Math. Theor. Comput. Sci.*, 22(4), 2020.
- [9] Z. Dvorák, D. Král, and R. Thomas. Testing first-order properties for subclasses of sparse graphs. *J. ACM*, 60(5):36:1–36:24, 2013.
- [10] K. Eickmeyer, A. C. Giannopoulou, S. Kreutzer, O. Kwon, M. Pilipczuk, R. Rabinovich, and S. Siebertz. Neighborhood complexity and kernelization for nowhere dense classes of graphs. In I. Chatzigiannakis, P. Indyk, F. Kuhn, and A. Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, volume 80 of *LIPICs*, pages 63:1–63:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
- [11] P. Erdős and R. Rado. Intersection theorems for systems of sets. *Journal of the London Mathematical Society*, s1-35(1):85–90, 1960.
- [12] G. Fabiański, M. Pilipczuk, S. Siebertz, and S. Toruńczyk. Progressive algorithms for domination and independence. In R. Niedermeier and C. Paul, editors, *36th International Symposium on Theoretical Aspects of Computer Science, STACS 2019, March 13-16, 2019, Berlin, Germany*, volume 126 of *LIPICs*, pages 27:1–27:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [13] J. Gajarský, M. Pilipczuk, and S. Toruńczyk. Stable graphs of bounded twin-width. In C. Baier and D. Fisman, editors, *LICS '22: 37th Annual ACM/IEEE Symposium on Logic in Computer Science, Haifa, Israel, August 2 - 5, 2022*, pages 39:1–39:12. ACM, 2022.
- [14] M. Grohe, S. Kreutzer, R. Rabinovich, S. Siebertz, and K. S. Stavropoulos. Coloring and covering nowhere dense graphs. *SIAM J. Discret. Math.*, 32(4):2467–2481, 2018.
- [15] M. Grohe, S. Kreutzer, and S. Siebertz. Deciding first-order properties of nowhere dense graphs. *J. ACM*, 64(3):17:1–17:32, 2017.
- [16] P. N. Klein. Multiple-source shortest paths in planar graphs. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2005, Vancouver, British Columbia, Canada, January 23-25, 2005*, pages 146–155. SIAM, 2005.

- [17] S. Kreutzer, R. Rabinovich, and S. Siebertz. Polynomial kernels and wideness properties of nowhere dense graph classes. *ACM Trans. Algorithms*, 15(2):24:1–24:19, 2019.
- [18] D. R. Lick and A. T. White. k -degenerate graphs. *Canadian Journal of Mathematics*, 22(5):1082–1096, 1970.
- [19] R. J. Lipton and R. E. Tarjan. Applications of a planar separator theorem. *SIAM J. Comput.*, 9(3):615–627, 1980.
- [20] M. Malliaris and S. Shelah. Regularity lemmas for stable graphs. *Transactions of the American Mathematical Society*, 366(3):1551–1585, 2014.
- [21] W. Nadara, M. Pilipczuk, R. Rabinovich, F. Reidl, and S. Siebertz. Empirical evaluation of approximation algorithms for generalized graph coloring and uniform quasi-wideness. *ACM J. Exp. Algorithmics*, 24(1):2.6:1–2.6:34, 2019.
- [22] J. Nederlof, M. Pilipczuk, and K. Węgrzycki. Bounding generalized coloring numbers of planar graphs using coin models. *CoRR*, abs/2201.09340, 2022.
- [23] J. Nešetřil and P. Ossona de Mendez. The grad of a graph and classes with bounded expansion. *Electron. Notes Discret. Math.*, 22:101–106, 2005.
- [24] J. Nešetřil and P. Ossona de Mendez. Linear time low tree-width partitions and algorithmic consequences. In J. M. Kleinberg, editor, *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, STOC 2006, Seattle, WA, USA, May 21–23, 2006*, pages 391–400. ACM, 2006.
- [25] J. Nešetřil and P. Ossona de Mendez. Tree-depth, subgraph coloring and homomorphism bounds. *Eur. J. Comb.*, 27(6):1022–1041, 2006.
- [26] J. Nešetřil and P. Ossona de Mendez. Grad and classes with bounded expansion I. Decompositions. *Eur. J. Comb.*, 29(3):760–776, 2008.
- [27] J. Nešetřil and P. Ossona de Mendez. First order properties on nowhere dense structures. *J. Symb. Log.*, 75(3):868–887, 2010.
- [28] J. Nešetřil and P. Ossona de Mendez. On nowhere dense graphs. *Eur. J. Comb.*, 32(4):600–617, 2011.
- [29] J. Nešetřil and P. Ossona de Mendez. *Sparsity - Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and Combinatorics*. Springer, 2012.
- [30] J. Nešetřil, P. Ossona de Mendez, M. Pilipczuk, R. Rabinovich, and S. Siebertz. Rankwidth meets stability. In D. Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021*, pages 2014–2033. SIAM, 2021.
- [31] J. Nešetřil, P. Ossona de Mendez, R. Rabinovich, and S. Siebertz. Classes of graphs with low complexity: The case of classes with bounded linear rankwidth. *Eur. J. Comb.*, 91:103223, 2021.
- [32] M. Pilipczuk, M. Pilipczuk, P. Sankowski, and E. J. van Leeuwen. Network sparsification for Steiner problems on planar and bounded-genus graphs. *ACM Trans. Algorithms*, 14(4):53:1–53:73, 2018.
- [33] M. Pilipczuk, M. Pilipczuk, and S. Siebertz. Lecture notes for the course “Sparsity” given at Faculty of Mathematics, Informatics, and Mechanics of the University of

Warsaw, Winter Semesters 2017/18 and 2019/20. Available at <https://www.mimuw.edu.pl/~mp248287/sparsity2>.

- [34] M. Pilipczuk and S. Siebertz. Polynomial bounds for centered colorings on proper minor-closed graph classes. *J. Comb. Theory, Ser. B*, 151:111–147, 2021.
- [35] M. Pilipczuk, S. Siebertz, and S. Toruńczyk. On the number of types in sparse graphs. In A. Dawar and E. Grädel, editors, *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, pages 799–808. ACM, 2018.
- [36] A. Pillay. *Introduction to stability theory*, volume 8 of *Oxford logic guides*. Clarendon Press, 1983.
- [37] K.-P. Podewski and M. Ziegler. Stable graphs. *Fundamenta Mathematicae*, 100(2):101–107, 1978.
- [38] F. Reidl, F. Sánchez Villaamil, and K. S. Stavropoulos. Characterising bounded expansion by neighbourhood complexity. *Eur. J. Comb.*, 75:152–168, 2019.
- [39] M. Sokołowski. Bounds on Half Graph Orders in Powers of Sparse Graphs. In J. Nešetřil, G. Perarnau, J. Rué, and O. Serra, editors, *Extended Abstracts EuroComb 2021*, pages 248–254, Cham, 2021. Springer International Publishing.
- [40] J. M. Steele. Variations on the monotone subsequence theme of Erdős and Szekeres. In D. Aldous, P. Diaconis, J. Spencer, and J. M. Steele, editors, *Discrete Probability and Algorithms*, pages 111–131, New York, NY, 1995. Springer New York.
- [41] G. Szekeres and H. S. Wilf. An inequality for the chromatic number of a graph. *Journal of Combinatorial Theory*, 4(1):1 – 3, 1968.
- [42] J. van den Heuvel, P. Ossona de Mendez, R. Rabinovich, and S. Siebertz. On the generalised colouring numbers of graphs that exclude a fixed minor. *Electron. Notes Discret. Math.*, 49:523–530, 2015.
- [43] D. West. *Introduction to Graph Theory*. Featured Titles for Graph Theory. Prentice Hall, 2001.
- [44] X. Zhu. Colouring graphs with bounded generalized colouring number. *Discret. Math.*, 309(18):5562–5568, 2009.