

INDUCED MINORS AND COARSE TREE DECOMPOSITIONS

MARIA CHUDNOVSKY[†], JULIEN CODSI[§], AJAYKRISHNAN E S[‡], AND DANIEL LOKSHTANOV[‡]

ABSTRACT. Let G be a graph, $S \subseteq V(G)$ be a vertex set in G and r be a positive integer. The *distance r -independence number* of S is the size of the largest subset $I \subseteq S$ such that no pair u, v of vertices in I has a path on at most r edges between them in G . It has been conjectured [Chudnovsky et al., arXiv, 2025] that for every positive integer t there exist positive integers c, d such that every graph G that excludes both the complete bipartite graph $K_{t,t}$ and the grid $\boxplus_t$ as an induced minor has a tree decomposition in which every bag has (distance 1) independence number at most $c(\log n)^d$. We prove a weaker version of this conjecture where every bag of the tree decomposition has distance $16(\log n + 1)$ -independence number at most $c(\log n)^d$. Along the way, we also prove a version of the conjecture where every bag of the decomposition has distance 8-independence number at most $2^{c(\log n)^{1-(1/d)}}$.

KEYWORDS: $K_{t,t}$ -induced-minor-free graphs, coarse tree decompositions, coarse Menger's theorem, distance r -independence number, tree-independence number.

1. INTRODUCTION

All graphs in this paper are finite and all logarithms are base 2. For standard graph theory terminology that is not defined here we refer the reader to Diestel's textbook [Die25]. Let $G = (V(G), E(G))$ be a graph. For a set $X \subseteq V(G)$, we denote by $G[X]$ the subgraph of G induced by X , and by $G - X$ the subgraph of G induced by $V(G) \setminus X$. For graphs G and H , we say that H is a *minor* of G , or that G *contains* H as a minor, if there exist disjoint connected induced subgraphs $\{X_v\}_{v \in V(H)}$ of G such that for every pair of distinct vertices u, v of H , if uv is an edge of H then there exists an edge from X_u to X_v in G . If, for every pair of distinct vertices u, v of H , uv is an edge of H if and only if there exists an edge from X_u to X_v in G , then H is an *induced minor* of G .

Given a family $\mathcal{H}$ of graphs, we say that a graph G is $\mathcal{H}$ -*induced-minor-free* if no induced minor of G is isomorphic to a member of $\mathcal{H}$. The *distance* between two vertices u and v in G is the minimum number of edges of a u - v path in G . An *independent* set in G is a set I of vertices of pairwise distance at least two. For a positive integer r , a *distance r -independent set* is a set I of vertices of pairwise distance strictly greater than r . Thus, a set is independent if and only if it is 1-independent. The *distance r -independence number* of a vertex set S is denoted by $\alpha_r(S)$ and defined as the size of the largest distance r -independent subset of S . The *independence number* of a set S is $\alpha(S) = \alpha_1(S)$.

For a graph G , a *tree decomposition* (T, χ) of G consists of a tree T and a map $\chi: V(T) \rightarrow 2^{V(G)}$ with the following properties:

- (1) For every $v \in V(G)$, there exists $t \in V(T)$ such that $v \in \chi(t)$.
- (2) For every $v_1 v_2 \in E(G)$, there exists $t \in V(T)$ such that $v_1, v_2 \in \chi(t)$.
- (3) For every $v \in V(G)$, the subgraph of T induced by $\{t \in V(T) \mid v \in \chi(t)\}$ is connected.

[†]Department of Mathematics, Princeton University, Princeton, NJ, USA. Supported by NSF Grants DMS-2348219 and CCF-2505100, AFOSR grant FA9550-25-1-0275, and a Guggenheim Fellowship.

[§]Department of Mathematics, Princeton University, Princeton, NJ, USA. Supported by NSF Grant DMS-2348219 and by the Fonds de recherche du Québec through the doctoral research scholarship 321124.

[‡]Department of Computer Science, University of California Santa Barbara, Santa Barbara, CA, USA Supported by NSF Grant CCF-2505099.

For each $t \in V(T)$, we refer to $\chi(t)$ as a *bag* of (T, χ) . The *width* of a tree decomposition (T, χ) , denoted by $\text{width}(T, \chi)$, is $\max_{t \in V(T)} |\chi(t)| - 1$. The *independence* number of a tree decomposition (T, χ) is the maximum independence number $\alpha(\chi(t))$ over all bags $\chi(t)$ for $t \in V(T)$. The *distance r -independence* number of a tree decomposition (T, χ) is the maximum distance r -independence number $\alpha_r(\chi(t))$ over all bags $\chi(t)$ for $t \in V(T)$. The *treewidth*, *tree independence* and *distance r -tree independence* number of G , are the minimum width, independence number, and distance r -independence number of a tree decomposition of G , respectively. We will say that a *class* of graphs has a bounded graph parameter (such as treewidth or tree-independence) if there exists a constant c such that the parameter is upper bounded by c for every graph in the class.

Treewidth has been extensively studied, as graphs whose treewidth is upper bounded by a constant have favorable structural [Bod98, RS86] and algorithmic [CFK⁺15] properties. Some of the applications (such as polynomial time solvability of a number of graph problems) carry over to graphs of bounded tree-independence [DMŠ24, LMM⁺24]. This motivates a systematic study of which classes have bounded treewidth or tree-independence.

In terms of minors, treewidth is well understood: there exists a fixed constant c such that, if t is the largest integer such that G contains the t times t grid $\boxplus_t$ as a minor, then the treewidth of G is at least t [Die25] and at most $t^9(\log t)^c$ for some fixed constant c [CC16, CT21]. On the other hand, recent results suggest [ST21] (see also [ABVT25]) that there does not exist a “nice” characterization of bounded treewidth graphs in terms of forbidden induced subgraphs. For example, for every sufficiently large n there exists an n -vertex graph which excludes a triangle K_3 as an induced subgraph, $K_{3,3}$ and $\boxplus_5$ as an induced minor (indeed even smaller graphs), and yet has treewidth (and also tree independence number) at least $c \cdot \log n$ for a constant $c > 0$ [ST21].

Quite often, problems that are solvable in polynomial time when the treewidth or tree-independence number of the input graph is a constant can be solved in *quasi-polynomial* ($O(n^{(\log n)^d})$ for some constant d) time when the treewidth or tree-independence number of the input graph is upper bounded by a poly-logarithmic (i.e. $O((\log n)^c)$ for some constant c) function in the number of vertices [LPR26]. Together with the above-mentioned difficulties of characterizing the graphs of bounded treewidth and tree independence in terms of induced minors, this motivates trying to characterize the graphs whose treewidth (or tree independence) is bounded poly-logarithmically in the number of vertices (see for example [ACHS22, CGH⁺25]).

It was conjectured [CCL⁺25] that, for every integer t there exist integers c, d such that every graph G either contains $K_{t,t}$ or $\boxplus_t$ as an induced minor or has tree-independence at most $(c \log n)^d$. A subset of the authors recently showed [CEL25] that characterizing the graphs with poly-logarithmic tree-independence is essentially equivalent to characterizing the graphs with poly-logarithmic treewidth and no clique on $c \cdot \log n$ vertices for a sufficiently large constant c .

The main results of this paper are the following three theorems.

Theorem 1.1. *Let t be a positive integer and let G be a $K_{t,t}$ -induced-minor-free graph. Then there exists constants $c(t), d(t)$ such that G either has $\boxplus_t$ as an induced minor or has a tree decomposition (T, χ) such that for every $x \in V(T)$, the distance 16 $\log 4n$ -independence number of the bag $\chi(x)$ is at most $115000 \cdot \log^3 4n \cdot f \log 4f$, where $f = (1000000 \cdot d \cdot \log^5 4n)^c$.*

Theorem 1.2. *Let t be a positive integer and let G be a graph that excludes $K_{t,t}$ and $\boxplus_t$ as induced minors. Then there exists a positive integer $c(t)$ and $\epsilon(t) \in (0, 1]$, such that G has a tree decomposition (T, χ) such that for every $x \in V(T)$, the distance 8-independence number of $\chi(x)$ is at most $2^{c(t)} \log^{1-\epsilon(t)} n$.*

Our results fit into a recent line of work concerning “coarse graph theory” [GP25, NSS25a, GKL23, HNST24, AHJ⁺24, NSS25b]. In this field tree decompositions where each bag can be covered by “not too many” balls of “small” radius are of interest. Another line of inquiry within coarse graph theory is to understand what versions of the classic Menger’s Theorem [Men27] hold

true when disjoint paths are replaced by paths that are sufficiently far apart, and bounded size separators are replaced by separators that can be covered by “not too many” balls of “small” radius. Note that by Lemma 6.2 of [CCFL26], the existence of a tree decomposition as above implies the Menger theorem analogue. Theorems 1.1 and 1.3 fit into this framework if “small” means $O(\log n)$, and “not too many” means $O((\log n)^c)$ for some constant c , while in Theorem 1.2 “small” is $O(1)$ and “not too many” is $O(2^{(\log n)^{1-\epsilon}})$ for some positive constant ϵ .

Theorem 1.3. *Let t be a positive integer, and let G be a $K_{t,t}$ -induced-minor-free graph. Let $A, B \subseteq V(G)$, and let k be a positive integer. Then there exist constants $d(t)$ and $\mu(t)$ such that one of the following holds:*

- (1) G has an A – B separator $S \subseteq V(G)$ whose distance $16 \log 4n$ -independence number is at most $100k \cdot g \cdot \log^3 4n$, where $g := (24 \log^2 4n + 4d(t) + 1)^{\mu(t)} \cdot (2\mu(t) + 1)^{4\mu(t)^2 + 1}$.
- (2) G contains a collection of k vertex-disjoint A – B paths that are pairwise anticomplete.

Proof Outline. The proofs of all main theorems begin with a reduction from the class of all $K_{t,t}$ -induced-minor-free graphs to the subclass of such graphs with degeneracy bounded by a constant. To obtain this reduction, we show that every $K_{t,t}$ -induced-minor-free graph G admits a partition of $V(G)$ into a constant number of parts such that each connected component of every part is contained in a ball of radius four in G (this is Theorem 9.2). Contracting each such component yields an induced minor G' of G with bounded chromatic number, and hence bounded clique number. Since G' remains $K_{t,t}$ -induced-minor-free, by a theorem in [BBCD24] or [GH25], we conclude that G' has bounded degeneracy. By construction, separators consisting of a bounded number of balls, as well as collections of vertex-disjoint and pairwise anticomplete paths in G' , can be lifted to corresponding structures in G . Therefore, for the remaining two theorems, it suffices to prove them for the subclass of $K_{t,t}$ -induced-minor-free graphs whose degeneracy is bounded by a constant.

A number of recent manuscripts consider the problem of efficiently finding (or proving the existence of) separators that can be covered by a small number of sets from a prescribed family $\mathcal{F}$, by rounding the corresponding linear programming relaxation. In general, these linear programs may have large integrality gaps. However, positive results are known when $\mathcal{F}$ has additional structure – for example when the sets are cliques [KLS⁺24] or when they have poly-logarithmic independence number [CEL25], where poly-logarithmic upper bounds on the integrality gaps can be obtained. We identify another type of $\mathcal{F}$, that we call layered families, which is amenable to the techniques used in these papers.

Let G be a graph and Π be an ordered partition of $V(G)$. We define the layered family of G corresponding to Π as follows. First, remove all edges whose endpoints lie in the same part of Π . Next, orient each remaining edge from the vertex in the part with higher index to the vertex in the part with lower index. For each source, that is, each vertex of indegree zero in the resulting orientation, include in the family the set of all vertices reachable from it by a directed path. In graphs of bounded degeneracy, there exists an ordered partition Π for which the layered family satisfies the following useful properties: (1) the integrality gap of the linear programming relaxation for separators using sets of this family is at most poly-logarithmic, (2) for every vertex v in the graph, $N[v]$ can be covered by a constant number of sets from $\mathcal{F}$, and (3) every set in $\mathcal{F}$ is contained in a ball of radius $O(\log n)$. The main technical contribution of our work is to prove property (1). After this, we follow the established framework of rounding the linear programming relaxation, as follows.

For Theorem 1.1, when the optimum objective value is “low”, we can round the linear program and obtain balanced separators that can be covered by a poly-logarithmic number of sets from the family $\mathcal{F}$ and when the objective value is “high”, we use the linear program to guide us in sampling an induced subgraph H with treewidth which is large compared to its poly-logarithmic maximum degree. This forces H , and hence G , to contain $\boxplus_t$ as an induced minor, since graphs excluding

both $K_{t,t}$ and $\boxplus_t$ as induced minors have treewidth polynomially bounded in their maximum degree, which in H is only poly-logarithmic (and therefore small).

For Theorem 1.2, we apply a theorem from [CCFL26] which asserts that graphs that exclude $K_{t,t}$ and $\boxplus_t$ as induced minors and K_t as a subgraph have sub-polynomial treewidth.

Finally, for Theorem 1.3, a linear program with “low” objective value can be rounded to obtain A – B separators that can be covered by a poly-logarithmic number of sets from $\mathcal{F}$. When the value is high, we sample an induced subgraph H of poly-logarithmic maximum degree containing many vertex-disjoint A – B paths. As H remains $K_{t,t}$ -induced-minor-free, we can apply a theorem from [BHKM23] to obtain an edge partition Π in which every connected component induced by a part is a star. But now, as H admits $\mathcal{P}$, A and B cannot be separated using a small number of vertices, and as maximum degree is poly-logarithmic, this property also holds for the induced minor of H obtained by contracting connected components of a fixed part in Π . This observation helps us find an induced subgraph H' of H that preserves a reasonable fraction of the A – B paths in $\mathcal{P}$, but has constant maximum degree. But, as H' has constant degree and admits a large number of vertex-disjoint A – B paths, by a theorem from [GKL23], H' and consequently G also admit a lot of vertex-disjoint and pairwise anticomplete A – B paths.

2. PRELIMINARIES

When clear from context, we denote by n the number of vertices of the considered graph G . We use $[k]$ to denote the set $\{1, 2, \dots, k\}$ for a positive integer k . Unless the base is specified, logarithms are in base 2. For a set A and a subset $B \subseteq A$, we write $B \subsetneq A$ to denote that B is a proper subset of A . A *clique* in a graph G is a set of pairwise adjacent vertices. For a positive integer t , K_t denotes a clique on t vertices. A *complete bipartite graph* is a graph G with vertex set $A \cup B$ such that a pair of vertices (u, v) form an edge if and only if $\{u, v\} \cap A$ contains a single vertex. For a positive integer t , $K_{t,t}$ denotes a complete bipartite graph where A and B contain exactly t vertices each. For every positive integer t , a t times t grid denoted as $\boxplus_t$ is the graph with vertex set $[t] \times [t]$ where every pair $(i, j), (i', j')$ of vertices are adjacent if and only if $|i - i'| + |j - j'| = 1$.

For graphs G and H , we say that H is a *minor* of G , or that G *contains* H as a minor, if there exist disjoint connected induced subgraphs $\{X_v\}_{v \in V(H)}$ of G such that for every pair of distinct vertices u, v of H , if (u, v) is an edge of H then there exists an edge from X_u to X_v in G . If, for every pair of distinct vertices u, v of H , (u, v) is an edge of H if and only if there exists an edge from X_u to X_v in G , then H is an *induced minor* of G .

Given a vertex subset $X \subseteq V(G)$ the subgraph of G induced by X is the graph $G[X]$ with vertex set X and edge set $\{(u, v) \in E(G) \mid u, v \in X\}$. We use $G - X$ to denote the graph induced by $V(G) \setminus X$. We use $N(v)$ to denote the set of vertices that are adjacent to v and $N[v]$ for $N(v) \cup \{v\}$. Similarly, $N(X)$ denotes $\{u \in V(G) \setminus X \mid u \in N(v) \text{ for some } v \in X\}$ and $N[X]$ denotes $N(X) \cup X$. The *degree* of v is $|N(v)|$ and the *maximum degree* of G , denoted $\Delta(G)$, is $\max_{v \in V(G)} \deg(v)$. A *walk* in a graph is a sequence of vertices in which each vertex is adjacent to the next, while a *path* is a walk in which every vertex is distinct. The *length* of a path is equal to the number of vertices on it. Given vertex subsets $A, B \subseteq V(G)$, an A – B *walk* (resp., A – B *path*) is a walk (resp., path) starting in A and ending in B . Likewise, for vertices $u, v \in V(G)$, a u – v *walk* or u – v *path* starts at u and ends at v . A path P in G is *induced* if $G[P]$ has exactly $|V(P)| - 1$ edges.

For positive integers k and r , we say that a subset S of the vertices of graph G is (k, r) –*coverable* if there exists a vertex subset $C \subseteq V(G)$ of size at most k such that every vertex of S has a path on at most r vertices to at least one vertex of C . A *distance r -independent set* is a set I of vertices such that no pair of vertices in I has a path between them on at most r vertices. We remark that the two notions are closely related. Suppose first that $S \subseteq V(G)$ is such that the largest distance- r -independent set $I \subseteq S$ in G has size k . Then S is (k, r) –coverable in G . Indeed, if there exists a vertex $v \in S$ that is not within distance r of any vertex of I , then $I \cup \{v\}$ is again distance- r -independent in G , contradicting the maximality of I . Conversely, suppose that S

is (k, r) -coverable in G , and let $C \subseteq V(G)$ be a set of size at most k witnessing this. We claim that every distance $2r$ -independent set $I \subseteq S$ has size at most k . Indeed, if $|I| > k$, then by the pigeonhole principle there exist distinct vertices $u, v \in I$ that are both within distance at most r of the same vertex of C . But this contradicts the assumption that I is distance $2r$ -independent.

For a positive real number ϕ , an (X, ϕ) -balanced separator is a vertex subset S such that every connected component of $G - S$ contains at most $\phi \cdot |X|$ vertices from X . If A, B are vertex subsets, then an A - B separator is a vertex subset S such that every A - B path in G contains at least one vertex from S . Let $\mathcal{F}$ be a family of vertex subsets of G . A fractional (A, B) -separator using $\mathcal{F}$ in G is an assignment $\{x_F\}_{F \in \mathcal{F}}$ of non-negative real numbers to elements of $\mathcal{F}$ such that $\sum_{F \in \mathcal{F}} x_F$ is at least one for every A - B path P . A fractional cover of X using $\mathcal{F}$ is an assignment $\{x_F\}_{F \in \mathcal{F}}$ of non-negative real numbers to elements of $\mathcal{F}$ such that $\sum_{F \in \mathcal{F}} x_F \geq 1$ for every $v \in X$. The fractional cover number of X with respect to $\mathcal{F}$ denoted $\text{fcov}_{\mathcal{F}}(X)$ is the minimum of $\sum_{F \in \mathcal{F}} x_F$ over all fractional covers $\{x_F\}_{F \in \mathcal{F}}$ of X . A subset $\mathcal{F}'$ of $\mathcal{F}$ is a cover of X using $\mathcal{F}$ if every vertex of X is contained in at least one element of $\mathcal{F}'$. The cover number of X using $\mathcal{F}$, denoted by $\text{cov}_{\mathcal{F}}(X)$ is the size of the smallest subfamily $\mathcal{F}'$ of $\mathcal{F}$ such that every vertex of X is contained in some element of $\mathcal{F}'$.

A Bernoulli random variable with success probability p , is one which takes the value 1 with probability p , and 0 with probability $1 - p$. We recall a version of the Chernoff bound. A stronger version of this result, along with its full proof, is available in [MU17].

Proposition 2.1 ([HR90]; [MU17]). *Let $X_1, X_2, \dots, X_\ell$ be independent Bernoulli random variables, and let $X = \sum_{i \in [\ell]} X_i$. If $R \geq 6\mathbb{E}[X]$, then*

$$\mathbb{P}[X \geq R] \leq 2^{-R}.$$

3. LAYERED FAMILY

Let G be a graph, and let $\Pi := (L_1, \dots, L_k)$ be an ordered partition of $V(G)$. We define the layer graph $\vec{G}$ of (G, Π) as the directed graph with $V(\vec{G}) := V(G)$ and $(u, v) \in E(\vec{G})$ if and only if $\{u, v\} \in E(G)$ and $u \in L_j, v \in L_i$ for some $i < j$. For vertices $u, v \in V(G)$, we say that v is a descendant of u , and u is an ancestor of v , if there exists a directed path from u to v in $\vec{G}$. Note that a vertex u is always an ancestor of itself. Similarly, v is a child of u , and u is a parent of v , if (u, v) is an arc of $\vec{G}$.

The layered family $\mathcal{F}$ corresponding to (G, Π) is defined as

$$\{\{v \in V(G) \mid v \text{ is a descendant of } u\} \mid u \in V(G) \text{ has no parent}\}.$$

The vertex u used to define such a set F is called the center of F . A family $\mathcal{F}$ of vertex subsets of G is called a layered family of G if there exists an ordered partition Π of $V(G)$ such that $\mathcal{F}$ is the layered family corresponding to (G, Π) . If k is the smallest positive integer such that every $F \in \mathcal{F}$ intersects at most k parts of Π , then we say that $\mathcal{F}$ has thickness k . We use $\mathcal{C}_{\mathcal{F}}$ to denote the set of centers of elements of $\mathcal{F}$. If $u \in \mathcal{C}_{\mathcal{F}}$, then we use F_u to denote the element of $\mathcal{F}$ that has u as a center. Similarly, $\mathcal{F}_S$ denotes the subfamily $\{F_u \in \mathcal{F} \mid u \in S\}$, where $S \subseteq \mathcal{C}_{\mathcal{F}}$. It follows from the definition that $\mathcal{F}$ covers $V(G)$ and that the center of each set in $\mathcal{F}$ is unique to that set. The following lemma is immediate from the above definitions.

Lemma 3.1. *Let G be a graph and k be a positive integer. If $\mathcal{F}$ is a layered family of G with thickness at most k , then for every $F \in \mathcal{F}$ and every $u, v \in F$, there exists a path in $G[F]$ from u to v that contains at most $2k - 1$ vertices, all of which are ancestors of u or v .*

Proof. Let Π be the ordered partition of $V(G)$ such that $\mathcal{F}$ is the layered family corresponding to (G, Π) . Let $F \in \mathcal{F}$ be a set with center w , and let $u \in F$. By definition, there exists a directed path from w to u in the layer graph of (G, Π) , contained within F . Let $P = (v_1, \dots, v_\ell)$ be such a path,

where $v_1 = w$ and $v_\ell = u$. Each vertex v_i is a child of v_{i-1} , and a parent of a vertex must lie in a part of Π with a strictly larger index; hence the vertices of P all lie in distinct parts of Π . Since $\mathcal{F}$ has thickness at most k , the path P uses at most k parts of Π , and therefore $\ell \leq k$. Thus, P is a path from w to u with at most k vertices, all of which are ancestors of u .

Now let u, v be two vertices in F and let P_u and P_v be the paths in $G[F]$ from u and v to the center of F , of length at most k , consisting solely of ancestors of u and v , respectively. These two paths imply the existence of a u - v walk Q of length at most $2k - 1$ contained within the set of ancestors of u or v . Hence, the shortest u - v path P in $G[Q]$ has length at most $2k - 1$ and lies within the same set. This completes the proof of the lemma. $\square$

Now, for a vertex subset $S \subseteq V(G)$, we define its *upward* and *downward closures* by

$$\hat{S} := \{u \in V(G) \mid \exists v \in S \text{ such that } u \text{ is an ancestor of } v\},$$

$$\check{S} := \{u \in V(G) \mid \exists v \in S \text{ such that } u \text{ is a descendant of } v\}.$$

We say that a vertex set is *upward* (resp., *downward*) *closed* if the upward (resp., downward) closure of the set is itself. Observe that the elements of $\mathcal{F}$ are downward closed by definition. We now prove some lemmas on the interactions between layered families and downward closed sets that will be useful in the next sections.

Lemma 3.2. *Let G be a graph with a layered family $\mathcal{F}$ of thickness at most k . If $S \subseteq V(G)$ is a downward closed set, then $\mathcal{F}' := \{F \setminus S \mid F \in \mathcal{F} \text{ and } F \setminus S \neq \emptyset\}$ is a layered family of $G - S$ with thickness at most k .*

Proof. Let $\Pi := (L_1, \dots, L_k)$ be an ordered partition of $V(G)$ such that $\mathcal{F}$ is the layered family corresponding to (G, Π) , and set $\Pi' := (L_1 \setminus S, \dots, L_k \setminus S)$. Let $\mathcal{F}''$ denote the layered family of $(G - S, \Pi')$; we show $\mathcal{F}' = \mathcal{F}''$, which gives both conclusions of the lemma since each $F \in \mathcal{F}$ intersects at most k parts of Π and consequently $F \setminus S$ intersects at most k parts of Π' .

We first show that $\mathcal{C}_{\mathcal{F}''} = \mathcal{C}_{\mathcal{F}'}$. Let $w \in V(G) \setminus S$. We claim that w is a center of $\mathcal{F}''$ if and only if $w \in \mathcal{C}_{\mathcal{F}} \setminus S$. We focus on the “if” direction, since the “only if” direction is immediate: any $w \in \mathcal{C}_{\mathcal{F}}$ has no parent in the layer graph of (G, Π) , and in particular none in $(G - S, \Pi')$. Suppose then that $w \notin \mathcal{C}_{\mathcal{F}}$, so w has a parent w' in the layer graph of (G, Π) . If $w' \in S$, downward closure of S would force the descendant w into S , contradicting $w \in V(G) \setminus S$. Hence $w' \in V(G) \setminus S$, which implies that w' is a parent of w in $(G - S, \Pi')$. Combined with the observation that for $w \in \mathcal{C}_{\mathcal{F}}$, downward closure of S gives $F_w \setminus S \neq \emptyset$ if and only if $w \notin S$, we obtain $\mathcal{C}_{\mathcal{F}''} = \mathcal{C}_{\mathcal{F}} \setminus S = \mathcal{C}_{\mathcal{F}'}$.

Fix $w \in \mathcal{C}_{\mathcal{F}'}$ and let $F \in \mathcal{F}$ be the element with center w . Any directed path from w in the layer graph of $(G - S, \Pi')$ is also one in the layer graph of (G, Π) , so the element of $\mathcal{F}''$ with center w is contained in $F \setminus S$. Conversely, let u be a vertex of $F \setminus S$, and consider any directed path $(w, v_1, \dots, v_{\ell-1}, u)$ from w to u in the layer graph of (G, Π) . We claim that every v_i lies in $V(G) \setminus S$. Indeed, u is a descendant of v_i ; if $v_i \in S$, downward closure would force $u \in S$, contradicting $u \in F \setminus S$. Hence the entire path lies in $V(G) \setminus S$, and is therefore a directed path in the layer graph of $(G - S, \Pi')$, witnessing $u \in F''$. The two elements with center w therefore coincide. Thus $\mathcal{F}' = \mathcal{F}''$, which completes the proof of the lemma. $\square$

Lemma 3.3. *Let G be a graph with a layered family $\mathcal{F}$ of thickness at most k . If C is a connected component of G , then $\mathcal{F}' := \{F \cap C \mid F \in \mathcal{F} \text{ and } F \cap C \neq \emptyset\}$ is a layered family of $G[C]$ with thickness at most k .*

Proof. Let $S := V(G) \setminus C$. We claim that S is downward closed. Indeed, suppose $v \in S$ and let u be a descendant of v in the layer graph of (G, Π) , where Π is any ordered partition of $V(G)$ such that $\mathcal{F}$ is the layered family corresponding to (G, Π) . Then there exists a directed path from v to u whose underlying edges all lie in G , so u and v lie in the same connected component of G . As

$v \notin C$, neither does u , and hence $u \in S$. Applying Lemma 3.2 to S now yields the result, since $G - S = G[C]$ and $F \setminus S = F \cap C$ for every $F \in \mathcal{F}$. $\square$

Lemma 3.4. *Let G be a graph with a layered family $\mathcal{F}$, and let $S \subseteq \mathcal{C}_{\mathcal{F}}$. Let C be a connected component of $G - \check{S}$, and let $Z := (\mathcal{C}_{\mathcal{F}} \cap C)$. Then $C \cup \check{S} = \check{S} \cup Z$.*

Proof. Let $\Pi := (L_1, \dots, L_k)$ be an ordered partition of $V(G)$ such that $\mathcal{F}$ is the layered family corresponding to (G, Π) . Since $\check{S} \cup Z = \check{S} \cup \check{Z}$, it suffices to show that $\check{Z} \cup \check{S} = C \cup \check{S}$. We first observe that for every directed path $(v_1, \dots, v_\ell)$ in the layer graph of (G, Π) with $v_\ell \notin \check{S}$, the vertex set $\{v_1, \dots, v_\ell\}$ is contained in $V(G) \setminus \check{S}$, and hence in a single connected component of $G - \check{S}$. Indeed, if $v_i \in \check{S}$ for some $i \in [\ell]$, downward closure of $\check{S}$ would force $v_\ell \in \check{S}$, contradicting $v_\ell \notin \check{S}$.

Now, for the inclusion $\check{Z} \cup \check{S} \subseteq C \cup \check{S}$, let $u \in \check{Z} \setminus \check{S}$. By definition, u is a descendant of some $w \in Z$ in the layer graph of (G, Π) , and the directed path from w to u ends outside $\check{S}$. Hence by our observation, this path lies in the connected component of $G - \check{S}$ containing w , namely C , giving $u \in C$. For the reverse inclusion, let $u \in C$. Since $\mathcal{F}$ covers $V(G)$, there exists $w \in \mathcal{C}_{\mathcal{F}}$ with $u \in F_w$, i.e., u is a descendant of w in the layer graph of (G, Π) . As $u \notin \check{S}$, the directed path from w to u avoids $\check{S}$ and lies in the connected component of $G - \check{S}$ containing u , namely C . In particular $w \in C \cap \mathcal{C}_{\mathcal{F}}$ which implies $u \in \check{Z}$. This completes the proof of the lemma. $\square$

Lemma 3.5. *Let G be a graph with a layered family $\mathcal{F}$ of thickness at most k . If $C \subseteq \mathcal{C}_{\mathcal{F}}$, then $\mathcal{F}' := \{F_w \mid w \in C\}$ is a layered family of $G[\check{C}]$ with thickness at most k .*

Proof. Let $\Pi := (L_1, \dots, L_k)$ be an ordered partition of $V(G)$ such that $\mathcal{F}$ is the layered family corresponding to (G, Π) , and set $\Pi' := (L_1 \cap \check{C}, \dots, L_k \cap \check{C})$. Let $\mathcal{F}''$ denote the layered family of $(G[\check{C}], \Pi')$; we show $\mathcal{F}' = \mathcal{F}''$, which gives both conclusions of the lemma since each $F_w \in \mathcal{F}'$ already intersects at most k parts of Π , and hence at most k parts of Π' .

Let $u \in \check{C}$. We claim that u is a center of $\mathcal{F}''$ if and only if $u \in C$. We focus on the “if” direction, since the “only if” direction is immediate: any $u \in C$ is a center of $\mathcal{F}$, hence has no parent in the layer graph of (G, Π) , and in particular none in $(G[\check{C}], \Pi')$. Suppose then that $u \notin C$. By definition of $\check{C}$, there exists $w \in C$ such that u is a descendant of w in the layer graph of (G, Π) . As $u \neq w$, the predecessor u' of u on a directed path from w to u is also a descendant of w , so $u' \in F_w \subseteq \check{C}$, and u' is a parent of u in $(G[\check{C}], \Pi')$. Hence $\mathcal{C}_{\mathcal{F}''} = C$.

Now fix $w \in C$ and let F_w'' denote the element of $\mathcal{F}''$ with center w . Any descendant of w in the layer graph of $(G[\check{C}], \Pi')$ is a descendant in the layer graph of (G, Π) as well, so $F_w'' \subseteq F_w$. Conversely, let $u \in F_w$ and let P be a directed path from w to u in the layer graph of (G, Π) . Every vertex on P is a descendant of $w \in C$, hence lies in $\check{C}$, so P is a path in the layer graph of $(G[\check{C}], \Pi')$, witnessing $u \in F_w''$. Therefore $F_w = F_w''$, and $\mathcal{F}' = \mathcal{F}''$. $\square$

We remark that Lemmas 3.2 and 3.3 are used in Section 6 and Lemmas 3.4 and 3.5 are applied in Section 8. Finally, for vertices $u, v \in V(G)$, a u - v path P is said to be *upward minimal* if there does not exist a u - v path P' such that $\hat{P}' \subsetneq \hat{P}$, or $\hat{P}' = \hat{P}$ and $|V(P')| < |V(P)|$. A path is *upward non-minimal* if it is not upward minimal. We need the following property of upward minimal paths.

Lemma 3.6. *Let G be a graph and k be a positive integer. If $\mathcal{F}$ is a layered family of G with thickness at most k , then for every $F \in \mathcal{F}$ and every upward minimal path P in G , it holds that $|F \cap P| \leq 2k - 1$.*

Proof. Suppose for contradiction that there exists F and P such that $|F \cap P| \geq 2k$. Let w denote the center of F and let $P = (v_1, \dots, v_i, \dots, v_{i'}, \dots, v_p)$, where v_i and $v_{i'}$ are the first and last vertices of F along P . Let $P_{v_i, v_{i'}}$ be the v_i - $v_{i'}$ path whose existence is asserted by Lemma 3.1. We construct a v_1 - v_p walk Q by replacing the v_i - $v_{i'}$ subpath of P with $P_{v_i, v_{i'}}$. The walk Q is strictly shorter than P , since we removed at least $2k$ vertices and replaced them with at most $2k - 1$. Furthermore,

$\hat{Q} = \hat{P}_{v_1, v_i} \cup \hat{P}_{v_i, v_{i'}} \cup \hat{P}_{v_{i'}, v_p}$. Note that $\hat{P}_{v_1, v_i} \subseteq \hat{P}$ and $\hat{P}_{v_{i'}, v_p} \subseteq \hat{P}$ since P_{v_1, v_i} and $P_{v_{i'}, v_p}$ are subpaths of P . Moreover, $\hat{P}_{v_i, v_{i'}} \subseteq \hat{P}$ as well, because $P_{v_i, v_{i'}}$ consists entirely of ancestors of v_i or $v_{i'}$, and is therefore contained in $\hat{P}$. Consequently, $\hat{Q} \subseteq \hat{P}$. Hence, if P' is the shortest induced v_1 – v_p path in $G[Q]$, then $|V(P')| < |V(P)|$ and $\hat{P}' \subseteq \hat{P}$, contradicting the upward minimality of P . Therefore, every $F \in \mathcal{F}$ and every upward minimal path P satisfy $|F \cap P| \leq 2k - 1$. $\square$

4. ROUNDING A–B SEPARATOR LP

Let A and B be vertex subsets of a graph G . Let $\mathcal{P}$ denote the set of all induced A – B paths. The linear program corresponding to an A – B separator using the elements of $\mathcal{F}$ is defined as follows, with nonnegative real variables x_F for every $F \in \mathcal{F}$.

$$\begin{aligned} \text{Minimize : } & \sum_{F \in \mathcal{F}} x_F \\ \text{Subject To : } & \sum_{\substack{F \in \mathcal{F} \\ F \cap P \neq \emptyset}} x_F \geq 1 & \forall P \in \mathcal{P} \end{aligned} \tag{1}$$

The main theorem that we prove in this section is the following:

Theorem 4.1. *Let G be an n -vertex graph, k be a positive integer and $\mathcal{F}$ be a layered family of G with thickness at most k . Let $A, B \subseteq V(G)$ be vertex subsets of G and suppose $\{x_F\}_{F \in \mathcal{F}}$ is an optimal solution to LP 1 corresponding to $(G, \mathcal{F}, A, B)$. Then there exists an A – B separator S in G such that*

$$\text{fcov}_{\mathcal{F}}(S) \leq 8k \log(2n) \cdot \sum_{F \in \mathcal{F}} x_F.$$

We begin with the following lemma.

Lemma 4.2. *Let G be an n -vertex graph, $\mathcal{F}$ be a family of vertex subsets, $A, B \subseteq V(G)$ be vertex subsets and $\{x_F\}_{F \in \mathcal{F}}$ be an optimal solution to LP 1. Then, there exists another solution, $\{y_F\}_{F \in \mathcal{F}}$ that satisfies $1 \geq y_F \geq \frac{1}{|\mathcal{F}|}$ for every $F \in \mathcal{F}$ with $y_F > 0$, and $\sum_{F \in \mathcal{F}} y_F \leq 2 \cdot \sum_{F \in \mathcal{F}} x_F$.*

Proof. Define $\{y_F\}_{F \in \mathcal{F}}$ as follows,

$$y_F := \begin{cases} 0 & \text{if } x_F < \frac{1}{2|\mathcal{F}|}, \\ \min\{1, 2x_F\} & \text{otherwise.} \end{cases}$$

and observe that $\sum_{F \in \mathcal{F}} y_F \leq 2 \cdot \sum_{F \in \mathcal{F}} x_F$. Furthermore, by construction, we have $1 \geq y_F \geq \frac{1}{|\mathcal{F}|}$ for every $F \in \mathcal{F}$ with $y_F > 0$. Finally, for every A – B path P we have

$$\sum_{\substack{F \in \mathcal{F} \\ F \cap P \neq \emptyset}} y_F \geq \min\{1, \sum_{\substack{F \in \mathcal{F} \\ F \cap P \neq \emptyset \\ x_F \geq \frac{1}{2|\mathcal{F}|}}} 2x_F\} \geq \min\left\{1, 2\left(1 - |\mathcal{F}| \cdot \frac{1}{2|\mathcal{F}|}\right)\right\} = 1,$$

where the second to last transition uses that $\{x_F\}_{F \in \mathcal{F}}$ is a solution to LP 1 and that there are at most $|\mathcal{F}|$ elements in $\mathcal{F}$ of value at most $\frac{1}{2|\mathcal{F}|}$. $\square$

Lemma 4.3. *Let G be a graph, k be a positive integer and $\mathcal{F}$ be a layered family of G with thickness at most k . Let $A, B \subseteq V(G)$ be vertex subsets of G and suppose $\{y_F\}_{F \in \mathcal{F}}$ is a solution to LP 1 corresponding to $(G, \mathcal{F}, A, B)$ that satisfies $y_F \geq \frac{1}{n}$ whenever $y_F > 0$. Then, there exists an A – B separator S in G such that $\text{fcov}_{\mathcal{F}}(S)$ is at most $4k \cdot \log(2n) \cdot \sum_{F \in \mathcal{F}} y_F$.*

Proof. For every vertex $v \in V(G)$, let $d_v = \min \{ \sum_{F \in \mathcal{F}, F \cap P \neq \emptyset} y_F \mid P \text{ is an } A\text{-}v \text{ path} \}$, let $y_v = \sum_{F \in \mathcal{F}, F \ni v} y_F$, and let $I_v = (d_v - y_v, d_v]$. Furthermore $d_v \leq \sum_{F \in \mathcal{F}, F \cap Q \neq \emptyset} y_F$ for every $A\text{-}v$ walk Q , since there exists an $A\text{-}v$ path in $G[Q]$.

Now, for every $r \in (0, 1]$, we define $S_r = \{v \in V(G) \mid r \in I_v\}$.

Claim 4.3.1. S_r forms an $A\text{-}B$ separator for every $r \in (0, 1]$.

Proof of the Claim. Let $r \in (0, 1]$ and $P = (v_1, v_2, \dots, v_p)$ be an $A\text{-}B$ path in G . Let v_i be the first vertex on P such that $d_{v_i} \geq r$. Such a vertex exists, since $v_p \in B$ and therefore $d_{v_p} \geq 1$, as $\{y_F\}_{F \in \mathcal{F}}$ is a solution to LP 1. If $i = 1$, then since $v_1 \in A$, we have $d_{v_1} = y_{v_1}$ which implies $v_i = v_1 \in S_r$. Otherwise, consider v_{i-1} . As the $A\text{-}v_i$ walk Q obtained by appending v_i to the $A\text{-}v_{i-1}$ path that realizes $d_{v_{i-1}}$ satisfies $\sum_{F \in \mathcal{F}, F \cap Q \neq \emptyset} y_F \leq d_{v_{i-1}} + y_{v_i}$, we must have $d_{v_i} - y_{v_i} \leq d_{v_{i-1}} < r$. Consequently $r \in I_{v_i}$ which implies that $v_i \in S_r$, and hence S_r is an $A\text{-}B$ separator. $\triangleleft$

Consider the random process where an $r \in (0, 1]$ is sampled uniformly at random.

Claim 4.3.2. $\mathbb{E}_r[\text{fcov}_{\mathcal{F}}(S_r)] \leq 4k \cdot \log(2n) \cdot \sum_{F \in \mathcal{F}} y_F$.

Proof of the Claim. For each $F \in \mathcal{F}$ and $r \in (0, 1]$ let us define $m_F^r = 0$ if $F \cap S_r = \emptyset$ and $\max \left(\{1\} \cup \left\{ \frac{1}{y_v} \mid v \in F \cap S_r \right\} \right)$ otherwise. Note that this quantity is well defined, since if $y_v = 0$, then $I_v = \emptyset$, and hence $r \notin I_v$, which implies $v \notin S_r$. Moreover, m_F^r is at most n , because $y_F \geq \frac{1}{n}$ whenever $y_F \neq 0$, and consequently $y_v \geq \frac{1}{n}$ whenever $y_v \neq 0$. Also, let $\hat{y}_F^r = y_F \cdot m_F^r$. We next show that $\{\hat{y}_F^r\}_{F \in \mathcal{F}}$ forms a fractional cover of S_r for every $r \in (0, 1]$. Note that if $u \in S_r$, then

$$\sum_{\substack{F \in \mathcal{F} \\ F \ni u}} \hat{y}_F^r \geq \sum_{\substack{F \in \mathcal{F} \\ F \ni u}} y_F \cdot \max \left(\{1\} \cup \left\{ \frac{1}{y_v} \mid v \in F \cap S_r \right\} \right) \geq y_u \cdot \max \left\{ 1, \frac{1}{y_u} \right\} \geq 1.$$

Therefore,

$$\begin{aligned} \mathbb{E}_r[\text{fcov}_{\mathcal{F}}(S_r)] &\leq \mathbb{E}_r \left[\sum_{F \in \mathcal{F}} \hat{y}_F^r \right] = \sum_{F \in \mathcal{F}} y_F \cdot \mathbb{E}_r[m_F^r] \\ &\leq \sum_{F \in \mathcal{F}} y_F \left[\sum_{i=0}^{\infty} 2^{i+1} \cdot \mathbb{P}_r[2^i \leq m_F^r < 2^{i+1}] \right] \\ &\leq \sum_{F \in \mathcal{F}} y_F \left[\sum_{i=0}^{\lfloor \log n \rfloor} 2^{i+1} \cdot \mathbb{P}_r[2^i \leq m_F^r] \right]. \end{aligned} \tag{2}$$

Now, in order to bound $\mathbb{P}_r[2^i \leq m_F^r]$, for each $0 \leq i \leq \lfloor \log n \rfloor$ and $F \in \mathcal{F}$, we define $F_i := \{v \in F \mid y_v \leq 2^{-i}\}$. Let $u, v \in F_i$ and P be the $u\text{-}v$ path in $G[F]$ of length at most $2k - 1$ that uses only ancestors of u or v , whose existence is asserted by Lemma 3.1. Let $\bar{v}$ be a vertex of this path, and without loss of generality, let $\bar{v}$ be an ancestor of v . Since v belongs to every set in $\mathcal{F}$ that contains $\bar{v}$, we have $y_{\bar{v}} \leq y_v$. Hence, $\bar{v} \in F_i$ and consequently P is a path in $G[F_i]$. Now, we define $I_{F_i} := \bigcup_{v \in F_i} I_v$, let u be the vertex in F_i with the smallest value of $d_u - y_u$, and let v be the vertex in F_i with the largest value of d_v . Concatenating the $u\text{-}v$ path of length at most $2k - 1$ in $G[F_i]$ with an $A\text{-}u$ path that realizes d_u yields an $A\text{-}v$ walk. Since this $A\text{-}v$ walk contains an $A\text{-}v$ path, we conclude that $d_v \leq d_u + (2k - 1) \cdot 2^{-i}$. Thus,

$$\mathbb{P}_r[2^i \leq m_F^r] = \mathbb{P}_r[\exists v \in F_i : r \in I_v] = \mathbb{P}_r[r \in I_{F_i}] \leq d_v - (d_u - 2^{-i}) \leq 2k \cdot 2^{-i}.$$

Combining this with Equation 2, we conclude that

$$\mathbb{E}_r[\text{fcov}_{\mathcal{F}}(S_r)] \leq \sum_{F \in \mathcal{F}} y_F \left[\sum_{i=0}^{\lfloor \log n \rfloor} 2^{i+1} \cdot 2k \cdot 2^{-i} \right] \leq 4k \cdot \log(2n) \cdot \sum_{F \in \mathcal{F}} y_F.$$

◁

Now, as $\mathbb{E}_r[\text{fcov}_{\mathcal{F}}(S_r)] \leq 4k \cdot \log(2n) \cdot \sum_{F \in \mathcal{F}} y_F$ by Claim 4.3.2, there must exist some $r_o \in (0, 1]$ such that $\text{fcov}_{\mathcal{F}}(S_{r_o}) \leq 4k \log(2n) \cdot \sum_{F \in \mathcal{F}} y_F$. Furthermore, S_{r_o} is an A – B separator by Claim 4.3.1, which concludes our proof. $\square$

Now we are ready to prove Theorem 4.1.

Theorem 4.1. *Let G be an n -vertex graph, k be a positive integer and $\mathcal{F}$ be a layered family of G with thickness at most k . Let $A, B \subseteq V(G)$ be vertex subsets of G and suppose $\{x_F\}_{F \in \mathcal{F}}$ is an optimal solution to LP 1 corresponding to $(G, \mathcal{F}, A, B)$. Then there exists an A – B separator S in G such that*

$$\text{fcov}_{\mathcal{F}}(S) \leq 8k \log(2n) \cdot \sum_{F \in \mathcal{F}} x_F.$$

Proof. Since $\mathcal{F}$ is a layered family of G , we have $|\mathcal{F}| \leq n$. We invoke Lemma 4.2 on the 5-tuple $(G, \mathcal{F}, A, B, \{y_F\}_{F \in \mathcal{F}})$ to obtain $\{y_F\}_{F \in \mathcal{F}}$, another solution to LP 1 that satisfies $y_F \geq \frac{1}{|\mathcal{F}|} \geq \frac{1}{n}$ whenever $y_F > 0$, and $\sum_{F \in \mathcal{F}} y_F \leq 2 \cdot \sum_{F \in \mathcal{F}} x_F$. Since $(G, \mathcal{F}, A, B, \{y_F\}_{F \in \mathcal{F}})$ satisfies the requirements of Lemma 4.3, we are guaranteed the existence of an A – B separator S in G such that $\text{fcov}_{\mathcal{F}}(S) \leq 4k \cdot \log(2n) \cdot \sum_{F \in \mathcal{F}} y_F \leq 8k \cdot \log(2n) \cdot \sum_{F \in \mathcal{F}} x_F$, which proves the theorem. $\square$

5. USING THE DUAL OF A – B SEPARATOR LP

The main theorem that we prove in this section is the following:

Theorem 5.1. *Let G be a graph, k be a positive integer and $\mathcal{F}$ be a layered family of G with thickness at most k . Let $A, B \subseteq V(G)$ be vertex subsets of G . Then, for every $\ell \geq \frac{1}{6} \log(4|\mathcal{F}|)$, there exists a multiset $\mathcal{Q}$ of induced A – B paths in G of cardinality at least ℓf , such that*

- *For every $F \in \mathcal{F}$, the number of paths in $\mathcal{Q}$ that intersect F is at most $6(\ell + 1)$.*
- *Each path in $\mathcal{Q}$ intersects F at most $2k - 1$ times.*

where f denotes the optimum value of LP 1 corresponding to $(G, \mathcal{F}, A, B)$.

Proof. Since the theorem holds trivially if $f = 0$, let us assume that f is positive. Note that due to the constraint in LP 1, f has to be at least 1. Let $\mathcal{P}$ denote the set of all induced A – B paths in G . We describe the dual of LP 1 using non-negative real variables $\{y_P\}_{P \in \mathcal{P}}$.

$$\begin{aligned} \text{Maximize : } & \sum_{P \in \mathcal{P}} y_P \\ \text{Subject To : } & \sum_{\substack{P \in \mathcal{P} \\ P \cap F \neq \emptyset}} y_P \leq 1 & \forall F \in \mathcal{F} \end{aligned} \tag{3}$$

Observe that, by strong duality [Won83], the dual LP also has the same optimum value f as the primal.

Claim 5.1.1. *There exists an optimal solution to LP 3 where, for every $P \in \mathcal{P}$ we have $y_P > 0$ only if P is upward minimal.*

Proof of the Claim. Consider the optimal solution $\{y_P\}_{P \in \mathcal{P}}$ to LP 3 with the smallest number of variables that have a non-zero assignment to upward non-minimal paths. For the sake of contradiction, let P' be an upward non-minimal path with $y_{P'} > 0$ and let P_o be any upward minimal

path that witnesses its upward non-minimality, i.e., let P_o be upward minimal and satisfy $\hat{P}_o \subsetneq \hat{P}'$, or $\hat{P}_o = \hat{P}'$ and $|V(P_o)| < |V(P')|$. We claim that the assignment $\{y'_P\}_{P \in \mathcal{P}}$, defined by

$$y'_P = \begin{cases} y_{P_o} + y_{P'}, & \text{if } P = P_o, \\ 0, & \text{if } P = P', \\ y_P, & \text{otherwise,} \end{cases}$$

is a feasible solution. This contradicts the definition of $\{y_P\}_{P \in \mathcal{P}}$, since $\{y'_P\}_{P \in \mathcal{P}}$ attains the same (optimal) objective function value but has one fewer nonzero variable that corresponds to an upward non-minimal path.

To this end, observe that any constraint of the LP that does not involve y'_{P_o} remains satisfied, since the values of all other variables in $\{y'_P\}_{P \in \mathcal{P}}$ do not exceed those in $\{y_P\}_{P \in \mathcal{P}}$. Now consider a constraint that involves y'_{P_o} and let F be the element of $\mathcal{F}$ it corresponds to. If $P_o \cap F \neq \emptyset$, then $\hat{P}_o \cap F \neq \emptyset$, and consequently $\hat{P}' \cap F \neq \emptyset$ as well. Moreover, if $\hat{P}' \cap F \neq \emptyset$, then as elements of $\mathcal{F}$ are downward closed, F must contain a vertex of P' , and hence $F \cap P' \neq \emptyset$. Therefore, any constraint that contains y'_{P_o} also contains $y'_{P'}$, and since $y'_{P_o} + y'_{P'} = y_{P'} + y_{P_o}$, these constraints also remain satisfied. Thus, all constraints remain satisfied under $\{y'_P\}_{P \in \mathcal{P}}$, which concludes the proof of the claim. $\triangleleft$

Let $\{y_P\}_{P \in \mathcal{P}}$ be an optimal solution to LP 3, satisfying $y_P > 0$ only if P is upward minimal. Observe that $\mathcal{D} := \{y_P/f\}_{P \in \mathcal{P}}$ is a probability distribution over $\mathcal{P}$, as every y_P takes non-negative values and satisfy $\sum_{P \in \mathcal{P}} y_P = f$. Let $\mathcal{Q} := \{Q_i\}_{i \in [f \cdot \ell]}$ be $[f \cdot \ell]$ independent samples drawn from the distribution $\mathcal{D}$ over the set $\mathcal{P}$. Let $F \in \mathcal{F}$, and let χ_F denote the number of paths in $\mathcal{Q}$ that intersect F . For every $i \in [f \cdot \ell]$, we have that,

$$\mathbb{P}_{Q_i \sim \mathcal{D}}[Q_i \cap F \neq \emptyset] = \sum_{\substack{P \in \mathcal{P} \\ P \cap F \neq \emptyset}} \mathbb{P}_{Q_i \sim \mathcal{D}}[Q_i = P] = \sum_{\substack{P \in \mathcal{P} \\ P \cap F \neq \emptyset}} \frac{y_P}{f} \leq \frac{1}{f}.$$

By linearity of expectation, combined with the fact that $f \geq 1$, it follows that the expected value of χ_F is at most $\ell + 1$. Applying union bound over all $F \in \mathcal{F}$ and the Chernoff bound from Theorem 2.1 to χ_F for every $F \in \mathcal{F}$, we get:

$$\mathbb{P}\left[\max_{F \in \mathcal{F}} \{\chi_F\} < 6(\ell + 1)\right] \geq 1 - \sum_{F \in \mathcal{F}} \mathbb{P}[\chi_F \geq 6(\ell + 1)] \geq 1 - |\mathcal{F}| \cdot 2^{-6(\ell+1)} \geq \frac{3}{4}.$$

Hence, with probability at least $\frac{3}{4}$, the sampled multiset has cardinality at least $f \cdot \ell$, consists of induced A - B paths in G , and satisfies that for every $F \in \mathcal{F}$, at most $6(\ell + 1)$ of these paths intersect F . Therefore, there must exist some multiset $\mathcal{Q}$ with these properties. Finally, observe that $\mathcal{Q}$ can only contain upward minimal paths since $\mathcal{D}$ assigns non-zero probability only to them. This, combined with Lemma 3.6, implies the statement of the theorem. $\square$

6. ROUNDING BALANCED SEPARATOR LP

We begin this section by defining the *balanced separator LP*. Let G be a graph, $\mathcal{F}$ be a layered family of G , $\mathcal{C}_{\mathcal{F}}$ be the set of centers of elements in $\mathcal{F}$ and $X \subseteq \mathcal{C}_{\mathcal{F}}$. For each pair $u, v \in X$, let $\mathcal{P}_{u,v}$ denote the set of all induced paths from u to v in G . We describe the balanced separator LP corresponding to the instance $(G, X, \mathcal{F})$, using non-negative real variables x_F and $d_{u,v}$, defined for

every $F \in \mathcal{F}$ and $u, v \in X$.

$$\begin{aligned}
& \text{Minimize:} && \sum_{F \in \mathcal{F}} x_F && (4) \\
& \text{Subject to:} && \sum_{v \in X} d_{u,v} \geq \frac{|X|}{10} && \forall u \in X \\
& && d_{u,v} \leq \sum_{\substack{F \in \mathcal{F} \\ F \cap P \neq \emptyset}} x_F && \forall u, v \in X, P \in \mathcal{P}_{u,v} \\
& && d_{u,v} \leq 1 && \forall u, v \in X
\end{aligned}$$

The linear program above is feasible and bounded: setting $x_F = 1$ for every $F \in \mathcal{F}$ and $d_{u,v} = 1$ for every $u, v \in X$ is a feasible assignment, since $V(G) \subseteq \bigcup_{F \in \mathcal{F}} F$ ensures that every induced path from u to v intersects some $F \in \mathcal{F}$. We also remark that it is a variant of the one used in the Leighton–Rao rounding scheme [LR99], adapted to operate on vertex sets rather than individual vertices, as in [KLS⁺24, CEL25]. We are now ready to state the main theorem proved in this section:

Theorem 6.1. *Let G be a graph, k be a positive integer and $\mathcal{F}$ be a layered family of G with thickness at most k . Then for every subset $X \subseteq \mathcal{C}_{\mathcal{F}}$, where $\mathcal{C}_{\mathcal{F}}$ denotes the set of centers of the elements of $\mathcal{F}$, there exists an $(X, \frac{95}{100})$ -balanced separator S in G such that S is downward closed and*

$$\text{fcov}_{\mathcal{F}}(S) \leq 5000kf \log(2n) \log(f+4),$$

whenever f , which denotes the optimum value of the balanced separator LP corresponding to $(G, X, \mathcal{F})$, is positive.

For the remainder of this section, we fix a graph G , the layered family $\mathcal{F}$ and the set $X \subseteq \mathcal{C}_{\mathcal{F}}$. We also fix an optimal solution to LP 4 corresponding to $(G, X, \mathcal{F})$ of value $f > 0$. Let x_F and $d_{u,v}$, for every $F \in \mathcal{F}$ and $u, v \in X$, denote the values assigned to the corresponding variables in this solution.

For $U \subseteq V(G)$, we define $\mu(U) := \sum_{\substack{F \in \mathcal{F} \\ F \cap U \neq \emptyset}} x_F$. For every $v \in V(G)$, we define $x_v := \sum_{\substack{F \in \mathcal{F} \\ F \ni v}} x_F$. Also, for every $u, v \in V(G)$, we define $d(u, v) := \min \{ \sum_{\substack{F \in \mathcal{F} \\ F \cap P \neq \emptyset}} x_F \mid P \in \mathcal{P}_{u,v} \}$. Note that $d(\cdot, \cdot)$ is a symmetric function, and that $d(v, v) = x_v$. We now make some more observations about the function d .

Observation 6.2. *If $u, v \in X$, then $\min\{d(\bar{u}, v), 1\} \geq d_{u,v}$.*

Observation 6.3. *If $u, v, w \in V(G)$, then $d(u, w) \leq d(u, v) + d(v, w)$. In particular, if $(v, w) \in E(G)$ and $u \in V(G)$, then $d(u, w) \leq d(u, v) + x_w$.*

Observation 6.2 follows immediately from the LP constraints: $d_{u,v} \leq \sum_{\substack{F \in \mathcal{F} \\ F \cap P \neq \emptyset}} x_F$ and $d_{u,v} \leq 1$ for every $u, v \in X$ and $P \in \mathcal{P}_{u,v}$. Furthermore, note that $d(u, v) \leq \sum_{\substack{F \in \mathcal{F} \\ F \cap Q \neq \emptyset}} x_F$ for any u – v walk Q in G , as we can always find an induced u – v path in $G[Q]$. Hence, Observation 6.3 follows from the fact that if P_1 and P_2 are induced u – v and v – w paths realizing $d(u, v)$ and $d(v, w)$ respectively, then appending P_2 to P_1 yields a u – w walk Q satisfying $\sum_{\substack{F \in \mathcal{F} \\ F \cap Q \neq \emptyset}} x_F \leq d(u, v) + d(v, w)$. Similarly, if $(v, w) \in E(G)$, then appending w to the u – v path realizing $d(u, v)$ produces a u – w walk satisfying $\sum_{\substack{F \in \mathcal{F} \\ F \cap Q \neq \emptyset}} x_F \leq d(u, v) + x_w$.

Let $\varepsilon := \frac{1}{500 \log(f+4)}$, and for every $u \in V(G)$, $C \subseteq V(G)$ and positive real number r , define

$$B_C(u, r) := \{v \in C \mid d(u, v) \leq r\} \quad \text{and} \quad \delta_C(u, r) := B_C(u, r + 3\varepsilon) \setminus B_C(u, r + \varepsilon).$$

Let $\ell_{\max} = \lceil \log \frac{f+4}{\varepsilon} \rceil$ and define $r_i := 3i\varepsilon$, where i is a positive integer. Let $Z_0 := \{v \in V(G) \mid x_v \geq \frac{\varepsilon}{2k}\}$. We note that every $F \in \mathcal{F}$ that contains a vertex $u \in Z_0$ also contains every descendant of u . This implies that $x_v \geq x_u$ and consequently $v \in Z_0$ for every descendant v of u . Hence Z_0 is downward closed.

Lemma 6.4. *Let u and v be vertices in $V(G) \setminus Z_0$ that satisfy $d(u, v) \geq \varepsilon$. Then there does not exist $F \in \mathcal{F}$ that contains both u and v .*

Proof. For the sake of contradiction, suppose there exists $F \in \mathcal{F}$ that contains both u and v . Let P be the u - v path in $G[F]$ of length at most $2k-1$ that uses only ancestors of u or v , whose existence is asserted by Lemma 3.1. Let $\bar{v}$ be a vertex of this path and without loss of generality let $\bar{v}$ be an ancestor of v . Since v belongs to every set in $\mathcal{F}$ that contains $\bar{v}$, we have $x_{\bar{v}} \leq x_v \leq \max\{x_u, x_v\}$. Hence $\bar{v} \in V(G) \setminus Z_0$, and consequently P is a path in $G - Z_0$. Thus, we have

$$d(u, v) \leq \sum_{\substack{F \in \mathcal{F} \\ F \cap P \neq \emptyset}} x_F < (2k-1) \cdot \frac{\varepsilon}{2k} < \varepsilon,$$

which gives the desired contradiction. $\square$

Now we state and prove three properties in Lemma 6.5, 6.6 and 6.7 that we require to implement a version of the Leighton-Rao rounding scheme [LR99] and prove Theorem 6.1.

Lemma 6.5. *Let $Z \subseteq V(G)$ be a superset of Z_0 . Suppose that $G - Z$ contains a connected component C such that $|C \cap X| > \frac{95|X|}{100}$, and let $\bar{u} \in X \cap C$ and $\ell \in [\ell_{\max}]$. Then we have,*

$$|X \setminus B_C(\bar{u}, r_{\ell+1})| \geq \frac{5|X|}{100}.$$

Proof. Since $\ell \in [\ell_{\max}]$, we have

$$r_{\ell+1} \leq r_{\ell_{\max}+1} = 3\varepsilon(\lceil \log \frac{f+4}{\varepsilon} \rceil + 1) \leq \frac{3 \log(2000 \cdot (f+4) \cdot \log(f+4))}{500 \log(f+4)} \leq \frac{1}{20},$$

where the third inequality upper bounds $\lceil \log \frac{f+4}{\varepsilon} \rceil + 1$ by $\log \frac{4(f+4)}{\varepsilon}$ and substitutes the value of ε and the final inequality follows from the facts that for $x \geq 4$ we have $\log x \geq 2$ and $\frac{\log \log x}{\log x} \leq \frac{3}{4}$.

Now, assume, for the sake of contradiction, that $|X \setminus B_C(\bar{u}, r_{\ell+1})| < \frac{5|X|}{100}$. Then we obtain:

$$\begin{aligned} \sum_{v \in X} \min\{d(\bar{u}, v), 1\} &= \sum_{v \in X \cap B_C(\bar{u}, r_{\ell+1})} \min\{d(\bar{u}, v), 1\} + \sum_{v \in X \setminus B_C(\bar{u}, r_{\ell+1})} \min\{d(\bar{u}, v), 1\} \\ &\leq (|X| - |X \setminus B_C(\bar{u}, r_{\ell+1})|) \cdot r_{\ell+1} + |X \setminus B_C(\bar{u}, r_{\ell+1})| \cdot 1 \\ &< \frac{95|X| \cdot r_{\ell+1}}{100} + \frac{5|X| \cdot 1}{100} \\ &< \frac{|X|}{10}. \end{aligned}$$

Here the second inequality follows from the fact that $r_{\ell+1} < 1$ and consequently maximizing $(|X| - |X \setminus B_C(\bar{u}, r_{\ell+1})|) \cdot r_{\ell+1} + |X \setminus B_C(\bar{u}, r_{\ell+1})| \cdot 1$ is equivalent to maximizing $|X \setminus B_C(\bar{u}, r_{\ell+1})|$. But, Observation 6.2 implies the inequality $\sum_{v \in X} \min\{d(\bar{u}, v), 1\} \geq \sum_{v \in X} d_{\bar{u}, v}$ which contradicts the constraint $\sum_{v \in X} d_{\bar{u}, v} \geq \frac{|X|}{10}$ of LP 4. Thus, it follows that $|X \setminus B_C(\bar{u}, r_{\ell+1})| \geq \frac{5|X|}{100}$. $\square$

Lemma 6.6. *Let $Z \subseteq V(G)$ be a superset of Z_0 . Suppose that $G - Z$ contains a connected component C such that $|C \cap X| > \frac{95|X|}{100}$, and let $\bar{u} \in X \cap C$. Then there exists $\ell \in [\ell_{\max}]$ such that $\mu(\delta_C(\bar{u}, r_\ell)) \leq \mu(B_C(\bar{u}, r_\ell))$.*

Proof. For this, we need the following claims:

Claim 6.6.1. $\mu(B_C(\bar{u}, r_1)) \geq 2\varepsilon$.

Proof of the Claim. Consider a path P in $G[C]$ from $\bar{u}$ to a vertex w such that $w \in X \cap (C \setminus B_C(\bar{u}, r_{\ell_{\max}+1}))$. Note that since $|X \setminus B_C(\bar{u}, r_{\ell+1})| \geq \frac{5|X|}{100}$ by Lemma 6.5 and since $|X \setminus C| < \frac{5|X|}{100}$ by assumption, such a vertex exists. Let $P' := (\bar{u}, v_1, \dots, v_p)$ be the subpath of P such that the successor of v_p in P (say v_{p+1}) is the first vertex in P that does not belong to $B_C(\bar{u}, r_1)$. P' is well defined, since w does not belong to $B_C(\bar{u}, r_1)$ and is non-empty since $\bar{u} \in B_C(\bar{u}, r_1)$. Hence, by Observation 6.3, we get:

$$\mu(B_C(\bar{u}, r_1)) \geq \sum_{\substack{F \in \mathcal{F} \\ F \cap P' \neq \emptyset}} x_F \geq d(\bar{u}, v_p) \geq d(\bar{u}, v_{p+1}) - x_{v_{p+1}} \geq 3\varepsilon - \frac{\varepsilon}{2k} > 2\varepsilon.$$

◁

Claim 6.6.2. $\mu(B_C(\bar{u}, r_{\ell+1})) \geq \mu(B_C(\bar{u}, r_\ell)) + \mu(\delta_C(\bar{u}, r_\ell))$ for every $\ell \in [\ell_{\max}]$.

Proof of the Claim. Let $v \in \delta_C(\bar{u}, r_\ell)$, and let $F \in \mathcal{F}$ be a set containing v , with w denoting the center of F . Observe that $d(\bar{u}, v) > r_\ell + \varepsilon$ while $d(\bar{u}, v') \leq r_\ell$ for every $v' \in B_C(\bar{u}, r_\ell)$. Hence by Observation 6.3, we get $d(v', v) \geq \varepsilon$, and by Lemma 6.4 we conclude that there does not exist an $F \in \mathcal{F}$ that contains both v and v' . As every set $F \in \mathcal{F}$ that intersects $\delta_C(\bar{u}, r_\ell)$ has an empty intersection with $B_C(\bar{u}, r_\ell)$, it follows that no set in $\mathcal{F}$ contributes to both $\mu(\delta_C(\bar{u}, r_\ell))$ and $\mu(B_C(\bar{u}, r_\ell))$. Moreover, from the definitions it follows that every set $F \in \mathcal{F}$ contributing to $\mu(B_C(\bar{u}, r_\ell)) + \mu(\delta_C(\bar{u}, r_\ell))$ also contributes to $\mu(B_C(\bar{u}, r_{\ell+1}))$. These two observations together imply the claim. ◁

Now, combining Claims 6.6.1 and 6.6.2, we conclude that if $\mu(\delta_C(\bar{u}, r_\ell)) > \mu(B_C(\bar{u}, r_\ell))$ for every $\ell \in [\ell_{\max}]$, then $\mu(B_C(\bar{u}, r_{\ell_{\max}})) > 2^{\ell_{\max}-1} \cdot \mu(B_C(\bar{u}, r_1)) \geq f$, while $\mu(G) \leq f$, which yields a contradiction. Hence we conclude that there exists some $\ell \in [\ell_{\max}]$ for which we have $\mu(\delta_C(\bar{u}, r_\ell)) \leq \mu(B_C(\bar{u}, r_\ell))$, which proves Lemma 6.6. ◻

Lemma 6.7. Let $Z \subseteq V(G)$ be a downward closed superset of Z_0 . Suppose that $G - Z$ contains a connected component C such that $|C \cap X| > \frac{95|X|}{100}$, and let $\bar{u} \in X \cap C$ and $\ell \in [\ell_{\max}]$. If $A := B_C(\bar{u}, r_\ell + \varepsilon)$, $B := C \setminus B_C(\bar{u}, r_{\ell+1})$ and $\mathcal{F}' := \{F \cap C \mid F \in \mathcal{F} \text{ and } F \cap C \neq \emptyset\}$, then there exists a fractional (A, B) -separator using $\mathcal{F}'$ in $G[C]$ of weight at most $\frac{1}{\varepsilon} \cdot \mu(\delta_C(\bar{u}, r_\ell))$.

Proof. Consider the assignment $\{x'_{F'}\}_{F' \in \mathcal{F}'}$ defined as follows. Let $F \in \mathcal{F}$ and let $F' \in \mathcal{F}'$ be the set corresponding to $F \cap C$. If $F \cap \delta_C(\bar{u}, r_\ell)$ is nonempty, define $x'_{F'} := \frac{x_F}{\varepsilon}$; otherwise, define $x'_{F'} := 0$. We claim that $\{x'_{F'}\}_{F' \in \mathcal{F}'}$ is a fractional (A, B) -separator using $\mathcal{F}'$ in $G[C]$.

Let $P = (v_0, v_1, \dots, v_p, v_{p+1})$ be a minimal (under subpaths) A - B path in $G[C]$. We have that $P \cap A = \{v_0\}$ and $P \cap B = \{v_{p+1}\}$. Let $P^* = (v_1, \dots, v_p)$ be the interior of P . By definition, we have $V(P^*) \subseteq \delta_C(\bar{u}, r_\ell)$. Moreover, P^* is non-empty; otherwise, applying Observation 6.3 to $\{\bar{u}, v_1, v_{p+1}\}$ yields $d(\bar{u}, v_{p+1}) \leq d(\bar{u}, v_1) + x_{v_{p+1}} \leq r_\ell + \varepsilon + \frac{\varepsilon}{2k}$, which contradicts the assumption that $d(\bar{u}, v_{p+1}) > r_\ell + 3\varepsilon$.

$$\begin{aligned} \sum_{\substack{F' \in \mathcal{F}' \\ F' \cap P \neq \emptyset}} x'_{F'} &\geq \sum_{\substack{F \in \mathcal{F} \\ F \cap P^* \neq \emptyset}} \frac{x_F}{\varepsilon} \geq \frac{1}{\varepsilon} \cdot d(v_1, v_p) \\ &\geq \frac{1}{\varepsilon} \cdot (d(\bar{u}, v_p) - d(\bar{u}, v_1)) \\ &\geq \frac{1}{\varepsilon} \left((r_\ell + 3\varepsilon - \frac{\varepsilon}{2k}) - (r_\ell + \varepsilon + \frac{\varepsilon}{2k}) \right) \geq 1, \end{aligned}$$

where the first inequality follows from the fact that any set in $\mathcal{F}$ that intersects P^* must intersect $\delta_C(\bar{u}, r_\ell)$. The second line follows from applying Observation 6.3 to the triple $\{\bar{u}, v_1, v_p\}$. The last applies Observation 6.3 to the triples $\{\bar{u}, v_p, v_{p+1}\}$ and $\{\bar{u}, v_0, v_1\}$ and uses $v_1, v_p \notin Z_0$. The claimed

upper bound on the weight of the fractional separator follows from the fact that $\sum_{F \in \mathcal{F}'} x'_F = \frac{1}{\varepsilon} \cdot \mu(\delta_C(\bar{u}, r_\ell))$. $\square$

With this, we prove the key lemma used in the proof of Theorem 6.1.

Lemma 6.8. *Let $Z \subseteq V(G)$ be a downward closed superset of Z_0 . If $G - Z$ has a connected component C such that $|C \cap X| > \frac{95|X|}{100}$ then there exists a partition $A \cup S \cup B$ of C such that:*

- (1) S is downward closed.
- (2) S is an A - B separator in $G - Z$.
- (3) $B \subsetneq C$.
- (4) $|A \cap X| \leq \frac{95|X|}{100}$.
- (5) $\text{fcov}_{\mathcal{F}}(S) \leq 4000k \cdot \log(2n) \cdot \log(f + 4) \cdot (\mu(C) - \mu(B))$.

Proof. Observe that by Lemma 3.2, $\mathcal{F}'' := \{F \setminus Z \mid F \in \mathcal{F} \text{ and } F \setminus Z \neq \emptyset\}$ is a layered family of $G - Z$ with thickness at most k . Furthermore, by Lemma 3.3 applied to the component C in $G - Z$ and layered family $\mathcal{F}''$, we conclude that

$$\mathcal{F}' := \{F'' \cap C \mid F'' \in \mathcal{F}'' \text{ and } F'' \cap C \neq \emptyset\} = \{F \cap C \mid F \in \mathcal{F} \text{ and } F \cap C \neq \emptyset\},$$

is a layered family of $G[C]$ with thickness at most k . Let $\bar{u} \in C \cap X$. We apply Lemma 6.6 to obtain some $\ell \in [\ell_{\max}]$ such that $\mu(\delta_C(\bar{u}, r_\ell)) \leq \mu(B_C(\bar{u}, r_\ell))$. Now, let $A' := B_C(\bar{u}, r_\ell + \varepsilon)$ and $B' := C \setminus B_C(\bar{u}, r_{\ell+1})$. We also know, by Lemma 6.7, that there exists a fractional (A', B') -separator using $\mathcal{F}'$ in $G[C]$ of weight at most $\frac{1}{\varepsilon} \cdot \mu(\delta_C(\bar{u}, r_\ell))$. Therefore, we can apply Theorem 4.1 to the tuple $(G[C], A', B', \mathcal{F}')$ to obtain an A' - B' separator S' in $G[C]$.

Define S to be equal to S' , B to be the set of vertices in connected components of $G[C \setminus S]$ that have a non-empty intersection with B' , and let $A := C \setminus (B \cup S)$. Observe that, by definition, S is both an A - B' separator and an A' - B separator in $G[C]$, and consequently also in $G - Z$. This establishes two facts. First, we have $|A \cap X| \leq |X| - |B' \cap X|$. By Lemma 6.5, $|B' \cap X| \geq \frac{5|X|}{100}$, which implies $|A \cap X| \leq \frac{95|X|}{100}$. Second, we obtain that $B \subsetneq C$ as $\bar{u} \in A'$ and consequently $\bar{u} \in C \setminus B$. Now, we bound the fractional cover number of S using $\mathcal{F}$ as follows:

$$\text{fcov}_{\mathcal{F}}(S) \leq \text{fcov}_{\mathcal{F}'}(S') \leq 8k \cdot \log(2n) \cdot \frac{1}{\varepsilon} \mu(\delta_C(\bar{u}, r_\ell)) \leq 4000k \cdot \log(2n) \cdot \log(f + 4) \cdot \mu(B_C(\bar{u}, r_\ell)).$$

Here, the first inequality follows from two observations. First, any fractional cover of S' using $\mathcal{F}'$ also covers S : every set containing a vertex of S' also contains all of its descendants, which implies that if a vertex of S' is covered, then all of its descendants are as well. Second, any set covered by $\mathcal{F}'$ can also be covered by $\mathcal{F}$, since $\mathcal{F}'$ is the restriction of $\mathcal{F}$ to a subset of $V(G)$. The second inequality follows from a property of S guaranteed by Theorem 4.1, while the last follows from our choice of ℓ and definition of ε .

To conclude the proof, it remains to show that $\mu(B_C(\bar{u}, r_\ell)) \leq \mu(C) - \mu(B)$. To this end, let $F \in \mathcal{F}$ be a set intersecting B , and let $v \in F \cap B$. Since $v \in B$ implies $v \notin A'$, we have $d(\bar{u}, v) > r_\ell + \varepsilon$, while $d(\bar{u}, v') \leq r_\ell$ for every $v' \in B_C(\bar{u}, r_\ell)$. Hence, by Observation 6.3, it follows that $d(v', v) \geq \varepsilon$. By Lemma 6.4, no set $F \in \mathcal{F}$ can contain both v and v' . Therefore, no set in $\mathcal{F}$ contributes to both $\mu(B)$ and $\mu(B_C(\bar{u}, r_\ell))$. Consequently $\mu(B_C(\bar{u}, r_\ell)) \leq \mu(C) - \mu(B)$, which completes the proof. $\square$

Now we are ready to prove Theorem 6.1.

Theorem 6.1. *Let G be a graph, k be a positive integer and $\mathcal{F}$ be a layered family of G with thickness at most k . Then for every subset $X \subseteq \mathcal{C}_{\mathcal{F}}$, where $\mathcal{C}_{\mathcal{F}}$ denotes the set of centers of the elements of $\mathcal{F}$, there exists an $(X, \frac{95}{100})$ -balanced separator S in G such that S is downward closed and*

$$\text{fcov}_{\mathcal{F}}(S) \leq 5000kf \log(2n) \log(f + 4),$$

whenever f , which denotes the optimum value of the balanced separator LP corresponding to $(G, X, \mathcal{F})$, is positive.

Proof. We prove the following claim:

Claim 6.1.1. *Let $Z \subseteq V(G)$ be a downward closed superset of Z_0 , and let C be the connected component of $G - Z$ such that $|C \cap X| > \frac{95|X|}{100}$, or let $C = \emptyset$ if no such component exists. Then there exists an $(X, \frac{95}{100})$ -balanced separator S in G such that S is downward closed and $\text{fcov}_{\mathcal{F}}(S) \leq \text{fcov}_{\mathcal{F}}(Z) + 4000k \cdot \log(2n) \cdot \log(f + 4) \cdot \mu(C)$.*

Proof of the Claim. We proceed by induction on $|C|$. Firstly, if $C = \emptyset$, then we can let $S = Z$ and the lemma holds true. Otherwise, let $A' \cup S' \cup B'$ be the partition of C whose existence is implied by Lemma 6.8. Observe that $Z \cup S'$ is downward closed, as both Z and S' are. Now let us consider the graph $G - (Z \cup S')$ and distinguish two cases. First, if there exists a component C' of $G - (Z \cup S')$ such that $|C' \cap X| > \frac{95|X|}{100}$, then $C' \subseteq B'$ as S' is an A' - B' separator in $G - Z$ and $|A' \cap X| \leq \frac{95|X|}{100}$. Since $B' \subsetneq C$, we conclude that $C' \subsetneq C$. Second, no such component exists; hence we set $C' = \emptyset$, and again $C' \subsetneq C$. In either case, by applying the inductive hypothesis to the pair $(G, Z \cup S')$, we obtain an $(X, \frac{95}{100})$ -balanced separator S in G such that S is downward closed and

$$\begin{aligned} \text{fcov}_{\mathcal{F}}(S) &\leq \text{fcov}_{\mathcal{F}}(Z \cup S') + 4000k \cdot \log(2n) \cdot \log(f + 4) \cdot \mu(C') \\ &\leq \text{fcov}_{\mathcal{F}}(Z) + \text{fcov}_{\mathcal{F}}(S') + 4000k \cdot \log(2n) \cdot \log(f + 4) \cdot \mu(C') \\ &\leq \text{fcov}_{\mathcal{F}}(Z) + 4000k \cdot \log(2n) \cdot \log(f + 4) \cdot (\mu(C) - \mu(B') + \mu(C')) \\ &\leq \text{fcov}_{\mathcal{F}}(Z) + 4000k \cdot \log(2n) \cdot \log(f + 4) \cdot \mu(C) \end{aligned}$$

where the last transition follows from the facts that $C' \subseteq B'$ which implies $\mu(C') \leq \mu(B')$. $\triangleleft$

Now, we apply Claim 6.1.1 to the pair (G, Z_0) and obtain an $(X, \frac{95}{100})$ -balanced separator S in G that is downward closed. Furthermore, since $\sum_{F \in \mathcal{F}} x_F \geq \frac{\varepsilon}{2k}$ for every $v \in Z_0$, it follows that $\{x_F \cdot \frac{2k}{\varepsilon}\}_{F \in \mathcal{F}}$ is a fractional cover for Z_0 . Consequently $\text{fcov}_{\mathcal{F}}(Z_0) \leq \frac{2k}{\varepsilon} \cdot f$.

Thus,

$$\begin{aligned} \text{fcov}_{\mathcal{F}}(S) &\leq \text{fcov}_{\mathcal{F}}(Z_0) + 4000k \cdot \log(2n) \cdot \log(f + 4) \cdot \mu(C) \\ &\leq 1000k \cdot \log(f + 4) \cdot f + 4000k \cdot \log(2n) \cdot \log(f + 4) \cdot f \\ &\leq 5000k \cdot \log(2n) \cdot \log(f + 4) \cdot f. \end{aligned}$$

which concludes the proof of Theorem 6.1. $\square$

7. USING THE DUAL OF BALANCED SEPARATOR LP

Let G be a graph, let $\mathcal{F}$ be a layered family of G , and let $\mathcal{C}_{\mathcal{F}}$ be the set of centers of elements in $\mathcal{F}$. Let $X \subseteq \mathcal{C}_{\mathcal{F}}$ and for each pair $u, v \in X$, let $\mathcal{P}_{u,v}$ denote the set of all induced paths from u to v in G . We now describe the dual of LP 4, corresponding to the instance $(G, X, \mathcal{F})$. It uses non-negative real variables ρ_u , $\eta_{u,v}$, and $\gamma_{u,v,P}$, defined for every $u, v \in X$ and path $P \in \mathcal{P}_{u,v}$.

$$\begin{aligned} \text{Maximize : } & \frac{|X|}{10} \sum_{u \in X} \rho_u - \sum_{u,v \in X} \eta_{u,v} \\ \text{Subject To : } & \rho_u - \eta_{u,v} - \sum_{P \in \mathcal{P}_{u,v}} \gamma_{u,v,P} \leq 0 & \forall u, v \in X \\ & \sum_{\substack{u,v \in X \\ P \in \mathcal{P}_{u,v} \\ P \cap F \neq \emptyset}} \gamma_{u,v,P} \leq 1 & \forall F \in \mathcal{F} \end{aligned} \tag{5}$$

It is easy to verify that the above is indeed the dual of LP 4, which by strong duality [Won83] immediately implies that the linear program is feasible and has bounded optimal objective function

value. Having established this, we are now ready to state the main theorem proved in this section. We remark that the general framework of this section closely follows that of Section 7 in [CEL25]. In that work, whenever the optimum value of LP 5 is high, one can sub-sample a set of paths that induce a subgraph of large treewidth. The main difference in this section is that we show it is possible to sample *upward minimal* paths that induce a subgraph of large treewidth. The upward minimality additionally guarantees that, for every sampled path P and every set $F \in \mathcal{F}$, the intersection $P \cap F$ has cardinality at most $2k - 1$, which is a property we require for our main results.

Theorem 7.1. *Let G be a graph, k be a positive integer and $\mathcal{F}$ be a layered family of G with thickness at most k . Then for every subset $X \subseteq \mathcal{C}_{\mathcal{F}}$, where $\mathcal{C}_{\mathcal{F}}$ denotes the set of centers of the elements of $\mathcal{F}$, and positive integer ℓ such that $\ell \geq 7 \cdot (f \cdot \log n + |X| + 2)$, there exists an induced subgraph $H \subseteq G$ with $X \subseteq V(H)$ that satisfies the following properties:*

- For every $(X, \frac{1}{2})$ -balanced separator S in H , we have that $\text{cov}_{\mathcal{F}}(S) \geq f$.
- For every $F \in \mathcal{F}$, we have that $|F \cap V(H)| \leq 1 + \frac{3\ell}{5f} \cdot (2k - 1)$.

whenever f , which denotes the optimum value of the dual of the balanced separator LP corresponding to $(G, X, \mathcal{F})$, is positive.

For the remainder of this section, we fix a graph G , the layered family $\mathcal{F}$ and the set $X \subseteq \mathcal{C}_{\mathcal{F}}$. We begin with the following lemma:

Lemma 7.2. *There exists an optimal solution to LP 5 where, for every $u, v \in X$ and $P \in \mathcal{P}_{u,v}$, we have $\gamma_{u,v,P} > 0$ only if P is upward minimal.*

Proof. Let ρ_u , $\eta_{u,v}$, and $\gamma_{u,v,P}$, for every $u, v \in X$ and $P \in \mathcal{P}_{u,v}$, denote the values assigned to the corresponding variables in an optimal solution to LP 5 that, among all optimal solutions, achieves the minimum number of nonzero γ variables corresponding to upward non-minimal paths. If for some $u', v' \in X$, there exists an upward non-minimal $u'-v'$ path P' with $\gamma_{u',v',P'} > 0$, then let P_o be any upward minimal path that witnesses its upward non-minimality, i.e., let P_o be upward minimal and satisfy $\hat{P}_o \subsetneq \hat{P}'$, or $\hat{P}_o = \hat{P}'$ and $|V(P_o)| < |V(P')|$. We claim that the assignment obtained from the above by setting, for all $u, v \in X$, $\rho'_u = \rho_u$ and $\eta'_{u,v} = \eta_{u,v}$, and by modifying the $\gamma_{u,v,P}$ variables as follows:

$$\gamma'_{u,v,P} = \begin{cases} \gamma_{u,v,P_o} + \gamma_{u,v,P'}, & \text{if } u = u', v = v', P = P_o, \\ 0, & \text{if } u = u', v = v', P = P', \\ \gamma_{u,v,P}, & \text{otherwise,} \end{cases}$$

is a feasible solution. This leads to a contradiction as the new solution attains the same (optimal) objective function value but has one fewer non-zero variable corresponding to an upward non-minimal path.

To this end, let us consider the constraint in LP 5. For ease of presentation, we refer to the constraints corresponding to vertex pairs in X as *type-one* constraints, and those corresponding to sets in $\mathcal{F}$ as *type-two* constraints. Any constraint of type-one that does not correspond to the pair u' and v' contains neither γ'_{u',v',P_o} nor $\gamma'_{u',v',P'}$ and therefore remains unchanged. The constraint of type-one corresponding to u' and v' contains both γ'_{u',v',P_o} and $\gamma'_{u',v',P'}$, and since $\gamma'_{u',v',P_o} + \gamma'_{u',v',P'} = \gamma_{u',v',P_o} + \gamma_{u',v',P'}$ its left-hand side remains unchanged. Consequently, this constraint also remains satisfied, which completes the case of type-one constraints. Now considering type-two, we note that any type-two constraint that does not involve γ'_{u',v',P_o} remains satisfied, since the values of all other γ' variables do not exceed those of γ variables. So consider a constraint that corresponds to element F in $\mathcal{F}$ that involves γ'_{u',v',P_o} , i.e., $P_o \cap F \neq \emptyset$. This implies that $\hat{P}_o \cap F$ is non-empty, and consequently $\hat{P}' \cap F$ is non-empty as well. Moreover, if $\hat{P}' \cap F \neq \emptyset$, then by definition of upward closure there exists a vertex of P' whose ancestor lies in F . But as elements of $\mathcal{F}$ are downward

closed, F contains a vertex of P' , and hence $F \cap P'$ is non-empty. Therefore, any constraint that contains γ'_{u',v',P_o} also contains $\gamma'_{u',v',P'}$, and since $\gamma'_{u',v',P_o} + \gamma'_{u',v',P'} = \gamma_{u',v',P_o} + \gamma_{u',v',P'}$ its left-hand side is unchanged. Consequently, this constraint remains satisfied as well, and therefore the new assignment is feasible, completing the proof of the lemma by contradiction. $\square$

Let us fix an optimal solution to LP 5, corresponding to $(G, X, \mathcal{F})$, that satisfies the conclusion of Lemma 7.2. Let ρ_u , $\eta_{u,v}$ and $\gamma_{u,v,P}$, for every $u, v \in X$ and $P \in \mathcal{P}_{u,v}$, denote the values assigned to the corresponding variables in this solution. Let f be the value of the objective function in this solution. For every $u, v \in X$, define $\gamma_{u,v} := \sum_{P \in \mathcal{P}_{u,v}} \gamma_{u,v,P}$, and let $\rho := \sum_{u \in X} \rho_u$.

Lemma 7.3. *The optimal solution has the following properties:*

- (1) $\sum_{v \in X} \eta_{u,v} \leq \frac{|X|}{10} \rho_u$ for every $u \in X$.
- (2) $10 \cdot f \leq \rho |X|$.

Proof. Given any feasible solution to the LP, we can always obtain a new assignment by choosing a u_0 in X and setting $\rho_{u_0} = 0$ and $\eta_{u_0,v} = 0$ for every $v \in X$. This preserves feasibility, since the only constraints involving these variables are of the form $\rho_{u_0} - \eta_{u_0,v} - \gamma_{u_0,v} \leq 0$, which remain satisfied even in the new assignment. Furthermore, the value of the objective function decreases by $\frac{|X|}{10} \rho_{u_0} - \sum_{v \in X} \eta_{u_0,v}$. But, if we start with the optimal solution, this modification should not strictly increase the objective function value. Therefore we get $\sum_{v \in X} \eta_{u_0,v} \leq \frac{|X|}{10} \rho_{u_0}$. Furthermore,

$$10f = 10 \left(\frac{|X|}{10} \sum_{u \in X} \rho_u - \sum_{u,v \in X} \eta_{u,v} \right) \leq 10 \left(\frac{|X|}{10} \sum_{u \in X} \rho_u \right) = \rho |X|.$$

$\square$

We define $\mathcal{P} = \bigcup_{u,v \in X} \mathcal{P}_{u,v}$, and introduce a probability distribution $\mathcal{D}$ over the domain $X \times X \times (\mathcal{P} \cup \{\emptyset\})$. The distribution $\mathcal{D}$ is defined using the following random process which has three steps: (Step 1) Select a vertex $u \in X$ with probability ρ_u/ρ . (Step 2) Select a vertex $v \in X$ uniformly at random. (Step 3) With probability $\eta_{u,v}/(\eta_{u,v} + \gamma_{u,v})$, output the triple $(u, v, \emptyset)$. Otherwise, choose a path $P \in \mathcal{P}_{u,v}$ with probability $\gamma_{u,v,P}/\gamma_{u,v}$ and output the triple (u, v, P) .

We now argue that the process is well defined. First, observe that $f > 0$ by assumption and $10 \cdot f \leq \rho |X|$ by Lemma 7.3, together imply $\rho > 0$. Hence $\{\rho_u/\rho\}_{u \in X}$ forms a probability distribution: each ρ_u is non-negative and $\sum_{u \in X} \rho_u = \rho$. Therefore, Step 1 is well defined. We next justify Step 3, as Step 2 is immediate. Let u be a vertex selected in Step 1. Then $\rho_u > 0$, which implies that $\eta_{u,v} + \gamma_{u,v} > 0$ for every $v \in X$, and in particular for the vertex v selected in Step 2; otherwise, the constraint $\rho_u - \eta_{u,v} - \gamma_{u,v} \leq 0$ of LP 5 would be violated. If $\gamma_{u,v} = 0$, then necessarily $\eta_{u,v}/(\eta_{u,v} + \gamma_{u,v}) = 1$, and the process outputs $(u, v, \emptyset)$. Thus, the second case of Step 3 occurs only when $\gamma_{u,v} > 0$. In this case, $\{\gamma_{u,v,P}/\gamma_{u,v}\}_{P \in \mathcal{P}_{u,v}}$ forms a probability distribution, since each $\gamma_{u,v,P}$ is non-negative and $\sum_{P \in \mathcal{P}_{u,v}} \gamma_{u,v,P} = \gamma_{u,v}$. Consequently, Step 3, and hence the entire process, is well defined.

Lemma 7.4. $\mathcal{D}$ satisfies $\mathbb{P}_{(u,v,P) \sim \mathcal{D}} [P = \emptyset] \leq \frac{1}{10}$.

Proof. Using Lemma 7.3, together with the fact that the constraint $\rho_u - \eta_{u,v} - \gamma_{u,v} \leq 0$ holds for every $u, v \in X$, we obtain:

$$\begin{aligned} \mathbb{P}_{(u,v,P) \sim \mathcal{D}} [P = \emptyset] &= \sum_{\substack{u' \in X \\ \rho_{u'} > 0}} \sum_{v' \in X} \mathbb{P}_{(u,v,P) \sim \mathcal{D}} [P = \emptyset \mid u = u', v = v'] \cdot \mathbb{P}_{(u,v,P) \sim \mathcal{D}} [v = v'] \cdot \mathbb{P}_{(u,v,P) \sim \mathcal{D}} [u = u'] \\ &\leq \frac{1}{|X|} \sum_{\substack{u' \in X \\ \rho_{u'} > 0}} \frac{\rho_{u'}}{\rho} \sum_{v' \in X} \frac{\eta_{u',v'}}{\eta_{u',v'} + \gamma_{u',v'}} \leq \frac{1}{|X|} \sum_{\substack{u' \in X \\ \rho_{u'} > 0}} \frac{1}{\rho} \cdot \frac{|X|}{10} \rho_{u'} = \frac{1}{10}. \end{aligned}$$

□

Lemma 7.5. *If $F \in \mathcal{F}$, then $\mathbb{P}_{(u,v,P) \sim \mathcal{D}} [P \cap F \neq \emptyset] \leq \frac{1}{10f}$.*

Proof. Let $u', v' \in X$ and $P' \in \mathcal{P}_{u',v'}$. Observe that if $\rho_{u'} = 0$ or $\gamma_{u',v'} = 0$, then it follows from the definition of the process that,

$$\mathbb{P}_{(u,v,P) \sim \mathcal{D}} [(u, v, P) = (u', v', P')] = 0.$$

Hence, let us consider the case where $\rho_{u'} > 0$ and $\gamma_{u',v'} > 0$. Using the fact that the constraint $\rho_u - \eta_{u,v} - \gamma_{u,v} \leq 0$ holds for every $u, v \in X$, we obtain:

$$\begin{aligned} \mathbb{P}_{(u,v,P) \sim \mathcal{D}} [(u, v, P) = (u', v', P')] &= \mathbb{P}_{(u,v,P) \sim \mathcal{D}} [P = P' \mid v = v', u = u'] \cdot \mathbb{P}_{(u,v,P) \sim \mathcal{D}} [v = v'] \cdot \mathbb{P}_{(u,v,P) \sim \mathcal{D}} [u = u'] \\ &= \frac{\gamma_{u',v',P'}}{\gamma_{u',v'}} \cdot \frac{\gamma_{u',v'}}{\gamma_{u',v'} + \eta_{u',v'}} \cdot \frac{1}{|X|} \cdot \frac{\rho_{u'}}{\rho} \leq \frac{\gamma_{u',v',P'}}{\rho|X|}. \end{aligned}$$

Now, let us fix $F \in \mathcal{F}$ and consider the probability that the sampled set P has a non-empty intersection with F .

$$\begin{aligned} \mathbb{P}_{(u,v,P) \sim \mathcal{D}} [P \cap F \neq \emptyset] &= \sum_{u',v' \in X} \sum_{\substack{P' \in \mathcal{P}_{u',v'} \\ P' \cap F \neq \emptyset}} \mathbb{P}_{(u,v,P) \sim \mathcal{D}} [(u, v, P) = (u', v', P')] \\ &\leq \frac{1}{\rho|X|} \sum_{\substack{u',v' \in X \\ P' \in \mathcal{P}_{u',v'} \\ \rho_{u'} > 0 \\ \gamma_{u',v'} > 0 \\ P' \cap F \neq \emptyset}} \gamma_{u',v',P'} \leq \frac{1}{\rho|X|} \leq \frac{1}{10f}. \end{aligned}$$

Here, the last inequality follows from Lemma 7.3, and the second-to-last follows from the fact that the constraint, $\sum_{\substack{u,v \in X \\ P \in \mathcal{P}_{u,v} \\ P \cap F \neq \emptyset}} \gamma_{u,v,P} \leq 1$, holds for every $F \in \mathcal{F}$. □

Finally, we require the following lemma, which relates balanced separators of the set A to A_1 - A_2 separators for a specific partition A_1, A_2 of A .

Lemma 7.6. *Let G be a graph and $A \subseteq V(G)$ be a vertex subset. If G has an $(A, \frac{1}{2})$ -balanced separator S , then there exists a partition A_1, A_2 of A such that $\max\{|A_1|, |A_2|\} \leq \frac{2|A|}{3}$ and S is an A_1 - A_2 separator.*

Proof. If $S \cap A$ contains at least $\frac{|A|}{2}$ elements, then selecting any subset $A_1 \subseteq S \cap A$ of size $\frac{|A|}{2}$ and setting $A_2 := A \setminus A_1$ suffices to prove the lemma. Otherwise, let $C_1, \dots, C_r$ be connected components of $G - S$ that have a non-empty intersection with A . Let $A'_i := A \cap C_i$ for every $i \in [r]$, and let $A'_{r+1} := A \cap S$. By relabeling the sets if necessary, assume that $|A'_i| \geq |A'_{i+1}|$ for all $i \in [r]$. Let q be the smallest integer in $[r+1]$ such that $\sum_{i=1}^q |A'_i| \geq \frac{|A|}{3}$. We claim that $\sum_{i=1}^q |A'_i| \leq \frac{2|A|}{3}$. If $q = 1$, the claim holds trivially, as S is an $(A, \frac{1}{2})$ -balanced separator and consequently $|A'_1| \leq \frac{|A|}{2}$. Otherwise, observe that

$$\sum_{i=1}^q |A'_i| = |A'_q| + \sum_{i=1}^{q-1} |A'_i| \leq |A'_{q-1}| + \sum_{i=1}^{q-1} |A'_i| \leq 2 \sum_{i=1}^{q-1} |A'_i| \leq \frac{2|A|}{3}.$$

Define $A_1 := \cup_{i=1}^q A'_i$ and $A_2 := A \setminus A_1$. Since the sets in $\{A'_i\}_{i \in [r+1]}$ are pairwise disjoint, we have that $|A_1| = \sum_{i=1}^q |A'_i|$. Hence, $|A_1| \leq \frac{2|A|}{3}$ and since $|A_1| \geq \frac{|A|}{3}$, we have $|A_2| \leq \frac{2|A|}{3}$. Furthermore, it follows from our construction of A_1, A_2 that S is an A_1 - A_2 separator, which concludes the proof of the lemma. □

Now we are ready to prove Theorem 7.1.

Theorem 7.1. *Let G be a graph, k be a positive integer and $\mathcal{F}$ be a layered family of G with thickness at most k . Then for every subset $X \subseteq \mathcal{C}_{\mathcal{F}}$, where $\mathcal{C}_{\mathcal{F}}$ denotes the set of centers of the elements of $\mathcal{F}$, and positive integer ℓ such that $\ell \geq 7 \cdot (f \cdot \log n + |X| + 2)$, there exists an induced subgraph $H \subseteq G$ with $X \subseteq V(H)$ that satisfies the following properties:*

- For every $(X, \frac{1}{2})$ -balanced separator S in H , we have that $\text{cov}_{\mathcal{F}}(S) \geq f$.
- For every $F \in \mathcal{F}$, we have that $|F \cap V(H)| \leq 1 + \frac{3\ell}{5f} \cdot (2k - 1)$.

whenever f , which denotes the optimum value of the dual of the balanced separator LP corresponding to $(G, X, \mathcal{F})$, is positive.

Proof. Let $\{(u_i, v_i, P_i)\}_{i \in [\ell]}$ be ℓ independent samples drawn from the distribution $\mathcal{D}$ over $X \times X \times (\mathcal{P} \cup \{\emptyset\})$. Let H be the subgraph of G induced by the union of all sampled sets and X , that is, $H = G[X \cup \bigcup_{i \in [\ell]} P_i]$. We claim that H has the desired properties with good probability.

Claim 7.6.1. *With probability at least $\frac{3}{4}$, the following holds: for every subfamily $\mathcal{F}' \subseteq \mathcal{F}$ of size at most f , and for every partition X_1, X_2 of X such that $\max\{|X_1|, |X_2|\} \leq \frac{2|X|}{3}$, there exists an X_1 - X_2 path in H that does not intersect any set in $\mathcal{F}'$.*

Proof of the Claim. We define an *eligible triple* to be a triple $(\mathcal{F}', X_1, X_2)$ such that $\mathcal{F}' \subseteq \mathcal{F}$ is a subfamily of size at most f , and X_1, X_2 is a partition of X satisfying $\max\{|X_1|, |X_2|\} \leq \frac{2|X|}{3}$. An eligible triple $(\mathcal{F}', X_1, X_2)$ is *bad* if every X_1 - X_2 path in H intersects at least one set in $\mathcal{F}'$. Let χ be the event that the conclusion of the claim is false, namely that there exists an eligible bad triple. To prove the claim it suffices to show $\mathbb{P}[\chi] \leq \frac{1}{4}$.

For every eligible triple $(\mathcal{F}', X_1, X_2)$, we define $\chi(\mathcal{F}', X_1, X_2)$ to be the event that this triple is bad. We will prove that for every eligible triple $(\mathcal{F}', X_1, X_2)$ we have $\mathbb{P}[\chi(\mathcal{F}', X_1, X_2)] \leq (\frac{9}{10})^\ell$, then $\mathbb{P}[\chi] \leq \frac{1}{4}$ follows by a simple union bound over all eligible triples. More concretely we have the following.

$$\mathbb{P}[\chi] \leq \left(\frac{9}{10}\right)^\ell \cdot 2^{|X|} \cdot \binom{|\mathcal{F}|}{|\mathcal{F}'|} \leq 2^{|X| + |\mathcal{F}'| \log |\mathcal{F}| - \ell \log(\frac{10}{9})} \leq \frac{1}{4}.$$

To prove that $\mathbb{P}[\chi(\mathcal{F}', X_1, X_2)] \leq (\frac{9}{10})^\ell$ we observe that the event $\chi(\mathcal{F}', X_1, X_2)$ does not occur if there exists an $i \in [\ell]$ such that P_i is an X_1 - X_2 path in H that is disjoint from every set in $\mathcal{F}'$. For every eligible triple $(\mathcal{F}', X_1, X_2)$ and every $i \in [\ell]$ we define the event $\bar{\chi}(\mathcal{F}', X_1, X_2, i)$ that P_i is an X_1 - X_2 path in H that is disjoint from every set in $\mathcal{F}'$. For every eligible triple $(\mathcal{F}', X_1, X_2)$ the events in $\{\bar{\chi}(\mathcal{F}', X_1, X_2, i) \mid i \in [\ell]\}$ are independent. Thus, to prove $\mathbb{P}[\chi(\mathcal{F}', X_1, X_2)] \leq (\frac{9}{10})^\ell$ it suffices to show that $\mathbb{P}[\bar{\chi}(\mathcal{F}', X_1, X_2, i)] \geq \frac{1}{10}$ for every $i \in [\ell]$.

To lower bound $\mathbb{P}[\bar{\chi}(\mathcal{F}', X_1, X_2, i)]$ we observe that (u_i, v_i, P_i) is sampled according to the distribution $\mathcal{D}$. The event $\bar{\chi}(\mathcal{F}', X_1, X_2, i)$ occurs unless $|\{u_i, v_i\} \cap X_1| \neq 1$, or $P_i = \emptyset$, or there exists $F \in \mathcal{F}'$ such that $P_i \cap F \neq \emptyset$. Since v_i is sampled uniformly from X we have that $\mathbb{P}[v_i \in X_1 | u_i \in X_1] \leq \frac{2}{3}$ and $\mathbb{P}[v_i \in X_2 | u_i \in X_2] \leq \frac{2}{3}$. The law of conditional probability applied to the event $u_i \in X_1$ now yields $\mathbb{P}[|\{u_i, v_i\} \cap X_1| \neq 1] \leq \frac{2}{3}$. By Lemma 7.4 we have that $\mathbb{P}[P_i = \emptyset] \leq \frac{1}{10}$. For every $F \in \mathcal{F}'$, by Lemma 7.5 we have that $\mathbb{P}[P_i \cap F \neq \emptyset] \leq \frac{1}{10f}$. A union bound over all $F \in \mathcal{F}'$ yields that the probability that there exists $F \in \mathcal{F}'$ such that $P_i \cap F \neq \emptyset$ is at most $\frac{|\mathcal{F}'|}{10f} \leq \frac{1}{10}$. We conclude that $\mathbb{P}[\bar{\chi}(\mathcal{F}', X_1, X_2, i)] \geq 1 - \frac{2}{3} - \frac{1}{10} - \frac{1}{10} \geq \frac{1}{10}$, giving the desired lower bound on $\mathbb{P}[\bar{\chi}(\mathcal{F}', X_1, X_2, i)]$, and completing the proof of the claim. $\triangleleft$

Claim 7.6.2. *With probability at least $\frac{3}{4}$, the following holds: For every $F \in \mathcal{F}$, we have that $|F \cap V(H)| \leq 1 + \frac{3\ell}{5f} \cdot (2k - 1)$.*

Proof of the Claim. Let $F \in \mathcal{F}$, and let χ_F denote the number of sets P_i , out of a total of ℓ samples, that intersect F . We define $\chi_{\mathcal{F}} := \max\{\chi_F \mid F \in \mathcal{F}\}$ and show that with probability at least $\frac{3}{4}$, the event $\chi_{\mathcal{F}} < 6 \cdot \frac{\ell}{10f}$ occurs.

Let $F \in \mathcal{F}$, since the probability that a fixed set P_i intersects F is at most $\frac{1}{10f}$, it follows from linearity of expectations that the expected value of χ_F is at most $\frac{\ell}{10f}$. Applying union bound over all $F \in \mathcal{F}$, the Chernoff bound from Theorem 2.1 to χ_F for every $F \in \mathcal{F}$, and using the lower bound on ℓ , we get:

$$\mathbb{P}\left[\chi_{\mathcal{F}} < 6 \cdot \frac{\ell}{10f}\right] \geq 1 - \sum_{F \in \mathcal{F}} \mathbb{P}\left[\chi_F \geq 6 \cdot \frac{\ell}{10f}\right] \geq 1 - |\mathcal{F}| \cdot 2^{-6 \cdot \frac{\ell}{10f}} \geq \frac{3}{4}.$$

Suppose that in our sampled subgraph H , the event $\chi_{\mathcal{F}} < \frac{3\ell}{5f}$ occurs. Since $\mathcal{D}$ only assigns non-zero probability to upward minimal paths, every sampled path is upward minimal. Hence, by Lemma 3.6, each path intersects a set in $\mathcal{F}$ at most $2k - 1$ times. Furthermore, since every set contains at most one element of X by definition, we get $|F \cap V(H)| \leq 1 + \frac{3\ell}{5f} \cdot (2k - 1)$. $\triangleleft$

Since the probability that the sampled subgraph satisfies both the properties stated in Claim 7.6.2 and Claim 7.6.1 is at least $\frac{1}{2}$, there exists an induced subgraph $H \subseteq G$ that satisfies both. Let H be such an induced subgraph. Combining Lemma 7.6 with Claim 7.6.1, we conclude that H has no $(X, \frac{1}{2})$ -balanced separator S with $\text{cov}_{\mathcal{F}}(S) \leq f$, thereby completing the proof of Theorem 7.1. $\square$

8. SOME RESULTS FOR AN AUXILIARY CLASS

Let $\mathcal{C}^c$ denote a family of graphs that is closed under the induced subgraph relation, with the additional property that every graph G in $\mathcal{C}^c$ satisfies $\text{tw}(G) \leq (\Delta(G))^c$. Let $\mathcal{C}_d^c$ be a subfamily of $\mathcal{C}^c$ which contains all graphs in $\mathcal{C}^c$ of degeneracy at most d . We begin with the following definition. Given a graph G and a positive integer d , a layered family is *d-neighborhood-witnessing* if for every vertex $v \in V(G)$, every $F \in \mathcal{F}$ that contains v satisfies $|N[v] \setminus F| \leq d$. The next lemma can be easily deduced from Theorem 7.1 of Abrishami, Chudnovsky, Hajebi and Spirkol [ACHS22], but we include a proof here for completeness.

Lemma 8.1. *Let d be a positive integer and G be a graph of degeneracy d . Then there exists a $4d$ -neighborhood-witnessing layered family $\mathcal{F}$ of G with thickness at most $\log(4n)$.*

Proof. We iteratively construct an ordered partition of $V(G)$ as follows. Starting with $i = 1$, define $V_i := V(G) \setminus \bigcup_{j=1}^{i-1} U_j$ and let U_i be the set of vertices in the induced subgraph $G[V_i]$ whose degree in $G[V_i]$ is at most $4d$. Let k be the last index for which V_k is non-empty. Let $\Pi := (U_1, U_2, \dots, U_k)$. Observe that Π is an ordered partition of $V(G)$ and let $\mathcal{F}$ be the layered family corresponding to (G, Π) .

First, we claim that k is at most $\log(4n)$ which directly implies that $\mathcal{F}$ can have thickness at most $\log(4n)$. For this claim, observe that any graph of degeneracy d has average degree at most $2d$, and consequently, at least half of its vertices have degree at most $4d$. Therefore, for every $i \in [k]$, the set U_i contains at least half of the vertices of V_i . Since $|V_1| = n$ and $|V_i| \leq \lfloor |V_{i-1}|/2 \rfloor$ for all $i \geq 2$, it follows that $|V_{\lceil \log n \rceil + 2}| = 0$ and by definition k is at most $\log(4n)$. For the other property, let v be a vertex in G , and $F \in \mathcal{F}$ be such that $v \in F$. Now, let U_i be the part in Π containing v , and observe that by definition, v has at most $4d$ neighbors in $V_i = \bigcup_{j \geq i} U_j$. Moreover, every neighbor of v in $V(G) \setminus V_i = \bigcup_{j < i} U_j$ is a descendant of v in the layer graph of (G, Π) . Therefore every $F \in \mathcal{F}$ that contains v satisfies $|N[v] \setminus F| \leq 4d$, which completes the proof of the lemma. $\square$

We note the following propositions, which we need for the proof of the main theorems.

Proposition 8.2 ([CFK⁺15]). *Let G be a graph of treewidth at most k , and let $X \subseteq V(G)$ be a vertex subset. Then there exists an $(X, \frac{1}{2})$ -balanced separator S in G of size at most $k + 1$.*

Proposition 8.3 ([Chv79]). *Let U be a set and let $\mathcal{F}$ be a family of subsets of U . If $S \subseteq U$ satisfies $\text{fcov}_{\mathcal{F}}(S) \leq k$, then it holds that $\text{cov}_{\mathcal{F}}(S) \leq k \cdot \log |U|$.*

With this, we are prepared to state and prove a key lemma which we require for the proof of Theorem 8.5.

Lemma 8.4. *Let c, d be positive integers and G be a graph in $\mathcal{C}_d^c$. Let k be a positive integer and $\mathcal{F}$ be a $4d$ -neighborhood-witnessing layered family of G with thickness at most k . Let $\mathcal{C}_{\mathcal{F}}$ be the set of centers of elements in $\mathcal{F}$, and $X \subseteq \mathcal{C}_{\mathcal{F}}$ be of size at most*

$$110000k \cdot \log^2(2n) \cdot f \log(4f); \quad \text{where } f = (1000000d \cdot k^2 \cdot \log^3(2n))^c.$$

Then there exists a set $S \subseteq \mathcal{C}_{\mathcal{F}}$ such that $\tilde{S} = \bigcup_{w \in S} F_w$ is an $(X, \frac{95}{100})$ -balanced separator in G and $|S| \leq 5000k \cdot \log^2(2n) \cdot f \log(4f)$.

Proof. Note that if $f \geq n$, then the lemma holds trivially by defining $S := \mathcal{C}_{\mathcal{F}}$, so let us assume that this is not the case. Consider the optimum value f' of LP 4 corresponding to $(G, X, \mathcal{F})$. If $f' \geq f$, then let $\ell := 7(f' \cdot \log n + 110000k \cdot \log^2(2n) \cdot f \log(4f) + 2)$ and apply Theorem 7.1 on $(G, \mathcal{F}, X, \ell)$ to obtain an induced subgraph $H \subseteq G$. Let v be a vertex of H and F be an element of $\mathcal{F}$ that contains v . By the properties of $\mathcal{F}$, we have $|N[v] \cap F| \leq 4d$. Moreover, since $|F \cap V(H)|$ is upper bounded by the properties of H asserted by Theorem 7.1, it follows that

$$\begin{aligned} |N[v] \cap V(H)| &\leq |(N[v] \cap V(H)) \cap F| + |(N[v] \cap V(H)) \setminus F| \\ &\leq 1 + \frac{3\ell}{5f'} \cdot (2k - 1) + 4d \\ &\leq 1 + \frac{21}{5} \cdot 2k \cdot (\log n + 110000k \cdot \log^2(2n) \cdot \log(4f) + 2) + 4d \\ &< (1 + \frac{42}{5} + \frac{42}{5} \cdot 110000 + \frac{84}{5} + 4) \cdot d \cdot k^2 \cdot \log^2(2n) \cdot \log(4f) \\ &< 1000000d \cdot k^2 \cdot \log^3(2n). \end{aligned}$$

Therefore, maximum degree of H is less than $1000000d \cdot k^2 \cdot \log^3(2n)$. Meanwhile, every $(X, \frac{1}{2})$ -balanced separator S in H satisfies $\text{cov}_{\mathcal{F}}(S) \geq f$, which implies that $|S| \geq f$. Applying the contrapositive of Proposition 8.2 to H , we conclude that $\text{tw}(H) \geq (1000000d \cdot k^2 \cdot \log^3(2n))^c$. But, these observations contradict the fact that since H is an induced subgraph of G , it must also belong to the family $\mathcal{C}_d^c$.

Hence f' is at most f , and consequently by Theorem 6.1, there exists an $(X, \frac{95}{100})$ -balanced separator S' in G satisfying $\text{fcov}_{\mathcal{F}}(S') \leq 5000k \cdot \log(2n) \cdot f \log(f + 4) \leq 5000k \cdot \log(2n) \cdot f \log(4f)$. Proposition 8.3 implies that $\text{cov}_{\mathcal{F}}(S') \leq 5000k \cdot \log^2(2n) \cdot f \log(4f)$. Thus, choosing any $S \subseteq \mathcal{C}_{\mathcal{F}}$ with $|S| \leq 5000k \cdot \log^2(2n) \cdot f \log(4f)$ such that $\mathcal{F}_S$ covers S' completes the proof of the lemma. $\square$

Theorem 8.5. *Let c, d be positive integers and G be a graph in $\mathcal{C}_d^c$. Let k be a positive integer and $\mathcal{F}$ be a $4d$ -neighborhood-witnessing layered family of G with thickness at most k . Let $\mathcal{C}_{\mathcal{F}}$ be the set of centers of elements in $\mathcal{F}$, and X be a subset of $\mathcal{C}_{\mathcal{F}}$. Then G has a rooted tree decomposition (T, χ) , with root $r \in V(T)$, such that $\chi(r) \supseteq \bigcup_{w \in X} F_w$ and for every $x \in V(T)$, there exists a subset $X' \subseteq \mathcal{C}_{\mathcal{F}}$ such that,*

$$|X'| \leq \max\{|X|, 110000k \cdot \log^2(2|V(G)|) \cdot f \log(4f)\} + 5000k \cdot \log^2(2|V(G)|) \cdot f \log(4f),$$

where $f = (1000000d \cdot k^2 \cdot \log^3(4|V(G)|))^c$; and the bag satisfies $\chi(x) = \bigcup_{w \in X'} F_w$.

Proof. We prove the theorem using induction on the size of $V(G)$. For the base case, when G has a single vertex, the theorem holds trivially. Assume inductively that the statement holds for all graphs with strictly fewer than n vertices, and consider a graph G on exactly n vertices. Note that if

$$|\mathcal{C}_{\mathcal{F}}| \leq \max\{|X|, 110000k \cdot \log^2(2n) \cdot f \log(4f)\} + 5000k \cdot \log^2(2n) \cdot f \log(4f),$$

then the theorem holds trivially, since a rooted tree decomposition with one node containing all of $V(G)$ satisfies the requirements. Hence, assume that this is not the case. Without loss of generality, also assume that $|X| \geq \lfloor 110000k \cdot \log^2(2n) \cdot f \log(4f) \rfloor$. Indeed, if this were not the case, we may add arbitrary elements of $\mathcal{C}_{\mathcal{F}}$ to X until equality holds, since proving the theorem for the resulting larger set immediately implies the statement for the original set. Now, if $|X| > 110000k \cdot \log^2(2n) \cdot f \log(4f)$, define $Y \subseteq X$ to be any subset of size $\lfloor 110000k \cdot \log^2(2n) \cdot f \log(4f) \rfloor$; otherwise, let $Y := X$. Let $S \subseteq \mathcal{C}_{\mathcal{F}}$ be the set obtained by applying Lemma 8.4 to the instance $(G, \mathcal{F}, Y)$. Recall that $\check{S} = \bigcup_{w \in S} F_w$, and let C be a connected component of $G - \check{S}$. Define $G_C := G[C \cup \check{S}]$, $\mathcal{F}_C := \{F_w \in \mathcal{F} \mid w \in S \cup (\mathcal{C}_{\mathcal{F}} \cap C)\}$, and $X_C := S \cup (X \cap C)$. We wish to apply the inductive assumption to the instance $(G_C, \mathcal{F}_C, X_C)$.

To justify this step, we need to establish two facts. First, let C be a connected component of $G - \check{S}$. Since $\check{S}$ is a $(Y, \frac{95}{100})$ -balanced separator, C contains at most $\frac{95}{100}|Y|$ vertices of Y . Since $|Y| = \lfloor 110000k \cdot \log^2(2n) \cdot f \log(4f) \rfloor$, at least $5500k \cdot \log^2(2n) \cdot f \log(4f)$ vertices of Y (and consequently of X since $Y \subseteq X$) must lie outside C . Finally, as $|S| \leq 5000k \cdot \log^2(2n) \cdot f \log(4f)$, at least $500k \cdot \log^2(2n) \cdot f \log(4f)$ vertices of X must lie outside G_C . Thus G_C has strictly fewer vertices than G , for every connected component C of $G - \check{S}$. Second, by Lemma 3.4, $C \cup \check{S}$ is the downward closure of $S \cup (\mathcal{C}_{\mathcal{F}} \cap C)$. Hence by Lemma 3.5, $\mathcal{F}_C$ is a layered family of G_C with thickness at most k .

Therefore it is safe to apply the theorem inductively to the instance $(G_C, \mathcal{F}'_C, X_C)$. Let (T_C, χ_C) denote the resulting rooted tree decomposition of G_C . We now define a rooted tree decomposition (T, χ) of G as follows. For the tree T we introduce a new root node r , and for each connected component C of $G - \check{S}$, connect r to the root of T_C . We define a bag function χ on $V(T)$ by setting $\chi(r) := \bigcup_{w \in X \cup S} F_w$, and for every other node $x \in V(T)$, letting $\chi(x) := \chi_C(x)$, where C is the connected component such that $x \in V(T_C)$. A straightforward verification of the definitions shows that the resulting pair (T, χ) is indeed a rooted tree decomposition of G . Furthermore, the construction satisfies $\chi(r) \supseteq \bigcup_{w \in X} F_w$. Now, let $x \in V(T)$. If $x = r$, then $\chi(x) = \bigcup_{w \in X \cup S} F_w$, and

$$|X \cup S| \leq |X| + 5000k \cdot \log^2(2|V(G)|) \cdot f \log(4f).$$

Otherwise, $x \in V(T_C)$ for some component C of $G - \check{S}$. Let $f_C := (1000000d \cdot k^2 \log^3(2|C \cup \check{S}|))^c$ and observe that by induction, there exists a set $X'_C \subseteq \mathcal{C}_{\mathcal{F}_C}$ such that $\chi(x) = \bigcup_{w \in X'_C} F_w$ and

$$\begin{aligned} |X'_C| &\leq \max\left\{|X_C|, 110000k \cdot \log^2(2|C \cup \check{S}|) \cdot f_C \log(4f_C)\right\} \\ &\quad + 5000k \cdot \log^2(2|C \cup \check{S}|) \cdot f_C \log(4f_C) \\ &\leq \max\left\{\left(|X| - \frac{5}{100}|Y| + 5000k \cdot \log^2(2n) \cdot f \log(4f)\right), 110000k \cdot \log^2(2n) \cdot f \log(4f)\right\} \\ &\quad + 5000k \cdot \log^2(2n) \cdot f \log(4f) \\ &\leq \max\left\{|X|, 110000k \cdot \log^2(2n) \cdot f \log(4f)\right\} \\ &\quad + 5000k \cdot \log^2(2n) \cdot f \log(4f). \end{aligned}$$

Since $\mathcal{C}_{\mathcal{F}_C} \subseteq \mathcal{C}_{\mathcal{F}}$, the same set X'_C witnesses the required property for the node x as an element of $V(T)$. Consequently, (T, χ) is a rooted tree decomposition satisfying all the desired properties, which completes the proof of the theorem. $\square$

Recall that for positive integers k and r , a subset $S \subseteq V(G)$ is said to be (k, r) -coverable if there exists a set $C \subseteq V(G)$ with $|C| \leq k$ such that every vertex of S has a path on at most r vertices to some vertex of C . Let c, d be positive integers, let $G \in \mathcal{C}_d^c$ be an n -vertex graph, and let $\mathcal{F}$ be a $4d$ -neighborhood-witnessing layered family in G with thickness at most $\log(4n)$. We apply the above theorem with $X = \emptyset$. Observe that if $F \in \mathcal{F}$ has center w , then F is contained in the ball

of radius $\log(4n)$ centered at w . Combining this observation with the theorem immediately yields the following.

Corollary 8.6. *Let c, d be positive integers and let G be a graph in $\mathcal{C}_d^c$. Then G has a tree decomposition (T, χ) such that for every $x \in V(T)$, the bag $\chi(x)$ is (k, r) -coverable in G , where k and r satisfy:*

$$\begin{aligned} r &\leq \log(4n). \\ k &\leq 115000 \cdot \log^3(4n) \cdot f \log(4f); \quad \text{where } f := (1000000d \cdot \log^5(4n))^c. \end{aligned}$$

Here we remark that the above corollary also follows from Lemma 8.4 combined with a theorem of Abrishami, Czyżewska, Kluk, Pilipczuk, Pilipczuk, and Rzażewski [ACK⁺25]. We also prove the following theorem regarding graphs of bounded degeneracy which we need towards the proof of Theorem 1.3.

Theorem 8.7. *Let d, f be positive integers and let G be an n -vertex graph of degeneracy d . Let $\mathcal{F}$ be a $4d$ -neighborhood-witnessing layered family of G with thickness at most $\lceil \log(2n) \rceil$, and let $A, B \subseteq V(G)$ be vertex subsets. Then, one of the following exists:*

- An A - B separator $S \subseteq V(G)$ in G , such that $\text{fcov}_{\mathcal{F}}(S) \leq 8 \log^2(4n) \cdot f$.
- An induced subgraph H of G , of maximum degree at most $24 \log^2(4n) + 4d$, such that every A_H - B_H separator S in H must have $\text{cov}_{\mathcal{F}}(S) \geq \frac{f}{12}$ in G , where $A_H := A \cap V(H)$ and $B_H := B \cap V(H)$.

Proof. Consider the optimum solution value f' of the LP 1 corresponding to $(G, \mathcal{F}, A, B)$. If $f' \leq f$, then applying Theorem 4.1 on $(G, \mathcal{F}, A, B)$ yields an A - B separator $S \subseteq V(G)$ in G , such that $\text{fcov}_{\mathcal{F}}(S) \leq 8 \log(4n) \cdot \log(2n) \cdot f' \leq 8 \log^2(4n) \cdot f$. Otherwise, we apply Theorem 5.1 to the instance $(G, \mathcal{F}, A, B)$ with $\ell := \log(2n)$, obtaining a multiset $\mathcal{Q}$ of A - B paths. Let H be the subgraph of G induced by the vertices that lie on paths in $\mathcal{Q}$. We claim that H has the desired properties. First, for every $F \in \mathcal{F}$, at most $6(\log(2n) + 1) \leq 12 \log(2n)$ paths of $\mathcal{Q}$ intersect F , while $\mathcal{Q}$ contains a total of $f \cdot \log(2n)$ paths. Therefore, any A_H - B_H separator S in H must satisfy $\text{cov}_{\mathcal{F}}(S) \geq \frac{f}{12}$. Second, let v be a vertex of H and let $F \in \mathcal{F}$. Since at most $12 \log(2n)$ paths of $\mathcal{Q}$ intersect F , and for every path $P \in \mathcal{Q}$ we have $|F \cap P| \leq 2 \log(4n) - 1$, it follows that $|F \cap V(H)| \leq 24 \log^2(4n)$. As $\mathcal{F}$ is $4d$ -neighborhood-witnessing, we conclude that $|N[v] \cap V(H)| \leq |F \cap V(H)| + 4d \leq 24 \log^2(4n) + 4d$, which completes the proof of the theorem. $\square$

9. VERTEX PARTITIONS OF GRAPHS WITH NO $K_{t,t}$ -INDUCED-MINORS

Let G be a graph, r a positive integer, and $S \subseteq V(G)$ a vertex subset. Recall that S is $(1, r)$ -coverable if there exists a vertex $v \in V(G)$ such that every vertex of S has a path on at most r vertices to v . A partition $(X_1, \dots, X_\ell)$ of S is called an r -radius vertex partition of S in G if for every $i \in [\ell]$ it holds that every connected component C of $G[X_i]$ is $(1, r)$ -coverable in G . Notice that if $u, v \in C$ for some $i \in [\ell]$ and some connected component C of $G[X_i]$, then there exists a u - v path on at most $2r$ vertices in G . We define $\pi_r(S, G)$ as the smallest integer ℓ for which there exists an r -radius vertex partition of S in G into ℓ parts. We simply write $\pi_r(G)$ to denote $\pi_r(V(G), G)$. We begin with the following lemma.

Lemma 9.1. *Let G be a graph and r be a positive integer. Then, for all $A, B \subseteq V(G)$,*

$$\max\{\pi_r(A, G), \pi_r(B, G)\} \leq \pi_r(A \cup B, G) \leq \pi_r(A, G) + \pi_r(B, G).$$

Furthermore, the first inequality holds with equality whenever A and B are anticomplete in G .

Proof. Set $p := \pi_r(A, G)$, $q := \pi_r(B, G)$, and let $(X_1, \dots, X_p)$, $(Y_1, \dots, Y_q)$ be optimal r -radius vertex partitions of A and B in G .

For the first inequality, let $(W_1, \dots, W_\ell)$ be an optimal r -radius vertex partition of $A \cup B$ in G . For each $j \in [\ell]$, every connected component of $G[W_j \cap A]$ is a subgraph of a connected component of $G[W_j]$, and hence $(1, r)$ -coverable. Thus $(W_1 \cap A, \dots, W_\ell \cap A)$ is an r -radius vertex partition of A , giving $\pi_r(A, G) \leq \ell$. By a symmetric argument, $\pi_r(B, G) \leq \ell$, and therefore $\max\{\pi_r(A, G), \pi_r(B, G)\} \leq \ell = \pi_r(A \cup B, G)$.

For the second inequality, define $Z_j := X_j$ for $j \in [p]$ and $Z_{p+j} := Y_j \setminus A$ for $j \in [q]$, and observe that $(Z_1, \dots, Z_{p+q})$ is a partition of $A \cup B$ into $p+q$ parts. Furthermore every connected component C of $G[Z_j]$ is a subgraph of a connected component C' of either $G[X_j]$ or $G[Y_{j-p}]$. Since C' is $(1, r)$ -coverable in G , the same center vertex witnesses $(1, r)$ -coverability of C . It follows that $(Z_1, \dots, Z_{p+q})$ is an r -radius vertex partition of $A \cup B$ in G , giving $\pi_r(A \cup B, G) \leq p+q$.

Now, suppose A and B are anticomplete. Set $\ell := \max\{p, q\}$ and, by appending empty parts if necessary, assume both partitions have exactly ℓ parts. Define $Z_j := X_j \cup Y_j$ for $j \in [\ell]$. Since $X_j \subseteq A$ and $Y_j \subseteq B$ are anticomplete, every connected component of $G[Z_j]$ lies entirely within $G[X_j]$ or entirely within $G[Y_j]$, and is therefore $(1, r)$ -coverable. Hence $\pi_r(A \cup B, G) \leq \ell = \max\{p, q\}$, which completes the proof of the lemma. $\square$

All of our main theorems rely on the following property of graphs that exclude $K_{t,t}$ as an induced minor:

Theorem 9.2. *Let G be a $K_{t,t}$ -induced-minor-free graph for a positive integer t . Then $\pi_4(G) \leq t \cdot 2^t$.*

For this, we need the following lemma.

Lemma 9.3. *Let G be a graph and let $Y \subseteq V(G)$ be such that $G[Y]$ is connected. Then there exist disjoint subsets T , M , and B of Y such that:*

- (1) $\pi_4(B, G) \geq \frac{\pi_4(Y, G)}{2}$.
- (2) $N(T) \cap Y \subseteq M$.
- (3) $G[T]$ and $G[B]$ are connected.
- (4) Every vertex $u \in B$ has a neighbor in M and every vertex $u \in M$ has a neighbor in T .

Proof. Let $\ell := \pi_4(Y, G)$. Within this proof, for every $u, v \in Y$, we define the *distance* between u and v as the length of the shortest u - v path in the induced subgraph $G[Y]$. Now fix an arbitrary vertex $v_0 \in Y$, and for every integer $i \geq 0$, let D_i denote the set of vertices in Y whose distance from v_0 is exactly i .

Claim 9.3.1. *There exists an index i such that $\pi_4(D_i, G) \geq \frac{\ell}{2}$.*

Proof of the Claim. Suppose for contradiction that $\pi_4(D_i, G) < \frac{\ell}{2}$ for every $i \geq 0$. Note that D_i and $D_{i'}$ are anticomplete in G whenever $|i - i'| \geq 2$. Indeed, any edge (v, v') with $v \in D_i$, $v' \in D_{i'}$, and $i' \geq i + 2$ yields a walk of length $i + 1$ from v_0 to v' in $G[Y]$, contradicting $v' \in D_{i'}$. Define $D_{\text{even}} := \bigcup_{i \geq 0} D_{2i}$ and $D_{\text{odd}} := \bigcup_{i \geq 0} D_{2i+1}$. Since any two even-indexed layers are anticomplete, the tightness condition of Lemma 9.1 gives $\pi_4(D_{\text{even}}, G) = \max\{\pi_4(D_{2i}, G) \mid i \geq 0\}$. The same argument applied to odd-indexed layers gives $\pi_4(D_{\text{odd}}, G) = \max\{\pi_4(D_{2i+1}, G) \mid i \geq 0\}$. The upper bound of Lemma 9.1 then yields $\pi_4(Y, G) \leq \pi_4(D_{\text{even}}, G) + \pi_4(D_{\text{odd}}, G) < \frac{\ell}{2} + \frac{\ell}{2} = \ell$, a contradiction. $\triangleleft$

Let i_o be such that $\pi_4(D_{i_o}, G) \geq \frac{\ell}{2}$, and let C be a component of $G[D_{i_o}]$ with $\pi_4(C, G) = \pi_4(D_{i_o}, G)$. Indeed such a component must exist: since distinct connected components of $G[D_{i_o}]$ are anticomplete by definition, the tightness condition of Lemma 9.1 implies $\pi_4(D_{i_o}, G) = \max_C \pi_4(C, G)$, where the maximum is over all components C of $G[D_{i_o}]$. Hence some component C satisfies $\pi_4(C, G) = \pi_4(D_{i_o}, G) \geq \frac{\ell}{2}$.

We claim that the sets $T := \bigcup_{i \leq i_o-2} D_i$, $M := D_{i_o-1}$, and $B := C$ satisfy all required conditions. Condition (1) holds because $B = C$ was chosen so that $\pi_4(C, G) \geq \frac{\ell}{2}$. For (2), suppose $v \in Y \setminus T$

has a neighbor $u \in T$. Since u lies at distance at most $i_o - 2$ from v_0 , we infer that v lies at distance at most $i_o - 1$, and hence $v \in M$, establishing $N(T) \cap Y \subseteq M$. Condition (3) follows from the fact that $G[B]$ is connected, as $B = C$ is a connected component of $G[D_{i_o}]$; and $G[T]$ is connected since every vertex in T has a path to v_0 entirely within T . Finally, for (4), observe that by construction every vertex in D_i (for $i \geq 1$) has a neighbor in D_{i-1} (its predecessor along a shortest v_0-v path in $G[Y]$), ensuring the necessary adjacency between T , M , and B . This concludes the proof of the lemma. $\square$

Now, we are ready to prove Theorem 9.2.

Theorem 9.2. *Let G be a $K_{t,t}$ -induced-minor-free graph for a positive integer t . Then $\pi_4(G) \leq t \cdot 2^t$.*

Proof. Suppose for contradiction that $\pi_4(G) > t \cdot 2^t$. We make the following claim.

Claim 9.2.1. *There exists a sequence of vertex subsets $V_1 \supset V_2 \supset \dots \supset V_{t+1}$ in G such that, for every $i \in [t]$, there exist disjoint subsets T_i and M_i of $V_i \setminus V_{i+1}$ such that:*

- (1) $G[T_i]$ is connected.
- (2) $N(T_i) \cap V_i \subseteq M_i$.
- (3) Every vertex $u \in M_i$ has a neighbor in T_i .
- (4) Every vertex $u \in V_{i+1}$ has a neighbor in M_i .

Moreover, there exists a subset $U \subseteq V_{t+1}$ of cardinality t such that for any distinct $u, v \in U$, there is no $u-v$ path of length at most four in G .

Proof of the Claim. Let $V_1 = V(G)$. Starting with $i = 1$, we iteratively construct the sets T_i , M_i , and B_i by applying Lemma 9.3 to V_i , and then define $V_{i+1} := B_i$. We repeat this process for each $i \in [t]$. This is valid because, by construction, each V_{i+1} is connected and therefore satisfies the conditions of the lemma in the next iteration. Observe that, by construction, $V_i \supset V_{i+1}$ for every $i \in [t]$. Moreover, the four required properties of the sets T_i , M_i , and V_{i+1} follow directly from the corresponding properties guaranteed by Lemma 9.3.

Let $U = \{u_1, \dots, u_{t'}\}$ be an inclusion-maximal subset such that no two distinct vertices of U are connected by a path of length at most four in G . We claim $t' \geq t$. For each $i \in [t']$, let X_i be the set of vertices of V_{t+1} that have a path of length at most four to u_i in G . Then $\pi_4(X_i, G) = 1$, since u_i witnesses (1,4)-coverability of X_i . Furthermore $(X_1, \dots, X_{t'})$ is a partition of V_{t+1} by maximality of U . Thus the upper bound of Lemma 9.1 gives $\pi_4(V_{t+1}, G) \leq t'$. But since the sets V_i satisfy the relation $\pi_4(V_{i+1}, G) \geq \frac{\pi_4(V_i, G)}{2}$, we conclude $\pi_4(V_{t+1}, G) \geq \frac{\pi_4(G)}{2^t} > t$, which completes the proof of the claim. $\triangleleft$

Let V_i , T_i and M_i for every $i \in [t]$, together with V_{t+1} and $U = \{u_1, \dots, u_t\} \subseteq V_{t+1}$, denote the sets whose existence is guaranteed by Claim 9.2.1. We construct an induced-minor model of $K_{t,t}$ in G , where sets $\{A_i\}_{i \in [t]}$ and $\{B_i\}_{i \in [t]}$ form the two sides of the bipartite graph in the model. For each $i \in [t]$, we define $A_i := T_i$ and $B_i := \{u_i\} \cup \{v_i^j\}_{j \in [t]}$, where v_i^j is a neighbor of u_i in M_j . Such a neighbor always exists because $u_i \in V_j$ for every $j \in [t]$ – this holds since $u_i \in V_{t+1}$ and, by construction, $V_{i'+1} \subseteq V_{i'}$ for all $i' \in [t]$ – and, by the properties of the construction, every vertex in V_{i+1} has at least one neighbor in M_i .

Now we need to argue that this is an induced-minor-model of $K_{t,t}$. Observe that by definition, A_i and B_i for all $i \in [t]$ induce connected subgraphs of G . Let $v \in A_i$, $v' \in A_{i'}$ for some $i > i'$. Note that $v \in V_i$ and consequently in $V_{i'+1}$, while $v' \in T_{i'}$. Therefore, v and v' are distinct as $V_{i'+1}$ and $T_{i'}$ are disjoint, and non-adjacent as $N(T_{i'}) \cap V_{i'}$ is contained within $M_{i'}$ which is disjoint from $V_{i'+1}$. Similarly, let $v \in B_i$, $v' \in B_{i'}$ for some $i \neq i'$. Note that if $(u_i, v, v', u_{i'})$ is a walk in G , it would contradict the fact that there is no $u_i-u_{i'}$ path on at most four vertices. Hence v and v' must be distinct and non-adjacent. Finally, let $i, j \in [t]$. We claim that there exists a vertex in A_i that is adjacent to $v_j^i \in B_j$. Indeed, by construction, v_j^i lies in M_i , and by Claim 9.2.1, every

vertex in M_i has a neighbor in T_i , which is precisely A_i . This proves that $\{A_i, B_i\}_{i \in [t]}$ forms an induced-minor-model of $K_{t,t}$ in G , contradicting our assumption that G is $K_{t,t}$ -induced-minor-free, and thereby completing the proof. $\square$

10. PROOF OF THEOREM 1.1

Towards the proof of Theorem 1.1, we state the following two propositions: the first is due to Hajebi, and the second due to Girão and Hunter; the latter also follows from a result of Bourneuf, Bucić, Cook, and Davies. We remark that the first proposition is stated in a weaker form than the result of Hajebi; however, this formulation suffices for our purposes.

Proposition 10.1 (Lemma 5.1 [Haj25]). *For every positive integer t , there exists a positive integer $c(t)$ such that any $\{K_{t,t}, \boxplus_t\}$ -induced-minor-free graph G satisfies $tw(G) \leq \Delta(G)^{c(t)}$.*

Proposition 10.2 (Theorem 1.5 [GH25], Theorem 1.1 [BBCD24]). *For every integer $k, t \geq 2$, there exists a positive integer $d(t, k)$ such that one of the following holds:*

- G contains K_k .
- G contains $K_{t,t}$ as an induced subgraph.
- G contains K_k as a proper induced subdivision.
- G has average degree at most $d(t, k)$.

Since any graph that contains a large clique as a proper induced subdivision also contains a large $K_{t,t}$ as an induced minor, and since forbidding a graph as a subgraph or as an induced minor is hereditary, Proposition 10.2 directly implies the following.

Proposition 10.3. *For every integer $t \geq 2$, there exists a positive integer $d(t)$ such that every graph that is K_t -free and $K_{t,t}$ -induced-minor-free has degeneracy at most $d(t)$*

With this, we are ready to prove Theorem 1.1.

Theorem 1.1. *Let t be a positive integer and let G be a $K_{t,t}$ -induced-minor-free graph. Then G either has $\boxplus_t$ as an induced minor or has a tree decomposition (T, χ) such that for every $x \in V(T)$, the distance $16 \log(4n)$ -independence number of the bag $\chi(x)$ is at most*

$$115000 \cdot \log^3(4n) \cdot f \log(4f); \quad \text{where } f = (1000000 \cdot d(t \cdot 2^t) \cdot \log^5(4n))^{c(t)}.$$

where c, d are the functions described in Propositions 10.1 and 10.3, respectively.

Proof. Assume that G has no $\boxplus_t$ -induced minor, as otherwise the statement holds. Consider a four-radius vertex partition $(X_1, \dots, X_{\pi_4(G)})$ of G with the minimum possible number of parts. We define G' to be the induced minor of G obtained by contracting, for each $i \in [\pi_4(G)]$ and every connected component C of $G[X_i]$, the edges of $G[C]$. Note that the vertex set of G' is $V(G') = \{v_C \mid C \text{ is a connected component of } G[X_i] \text{ for some } i \in [\pi_4(G)]\}$, and two vertices v_{C_1} and v_{C_2} are adjacent in G' whenever there exist vertices $u_1 \in C_1$ and $u_2 \in C_2$ with $(u_1, u_2) \in E(G)$.

Since G' is obtained by contracting edges of a $\{K_{t,t}, \boxplus_t\}$ -induced-minor-free graph, it is itself $\{K_{t,t}, \boxplus_t\}$ -induced-minor-free. Therefore, by Proposition 10.1, we have $tw(G') \leq \Delta(G')^{c(t)}$. Furthermore, for every $i \in [\pi_4(G)]$, the set $\{v_C \mid C \text{ is a connected component of } G[X_i]\}$ is an independent set in G' . Hence the chromatic number of G' , and consequently the size of its largest clique, is at most $\pi_4(G) \leq t \cdot 2^t$. By Proposition 10.3, it follows that G' has degeneracy at most $d \leq d(t \cdot 2^t)$. Consequently, $G' \in \mathcal{C}_d^c$ for $c := c(t)$ and $d := d(t \cdot 2^t)$. Applying Corollary 8.6 to G' we obtain a tree decomposition (T', χ') such that for every $x \in V(T')$, the bag $\chi'(x)$ is (k', r') -coverable in G' , for $k' \leq 115000 \cdot \log^3(4n) \cdot f \log(4f)$, and $r' \leq \log(4n)$ where f equals $(1000000d \cdot \log^5(4n))^c$.

We define a pair (T, χ) as follows. Let $T := T'$, and for every node $x \in V(T)$ define $\chi(x) := \{u \in V(G) \mid u \in C \subseteq X_i \text{ and } v_C \in \chi'(x)\}$. A direct application of definitions suffices to verify that (T, χ) is indeed a tree decomposition of G . Now, consider $x \in V(T)$ and let I be a largest

$16 \log(4n)$ -independent set of G contained in $\chi(x)$. Let $I' := \{v_C \in V(G') \mid C \cap I \neq \emptyset\}$. Observe that $|I'| = |I|$. Indeed, every vertex of I belongs to some connected component C of $G[X_i]$ for some $i \in [\pi_4(G)]$. Moreover, no two vertices $u, v \in I$ belong to the same connected component C of $G[X_i]$, since each such component is $(1, 4)$ -coverable and hence contains a u - v path on at most eight vertices, contradicting the $16 \log(4n)$ -independence of I . Let $C' \subseteq V(G')$ be a set witnessing that $\chi'(x)$ is (k', r') -coverable, i.e., $|C'| \leq 115000 \cdot \log^3(4n) \cdot f \log(4f)$ and for every vertex v_C of $\chi'(x)$, there exists at least one vertex in C' that has a path on at most $\log(4n)$ vertices to v_C . If $|I'| > 115000 \cdot \log^3(4n) \cdot f \log(4f)$, then by the pigeonhole principle there exists a vertex of C' that has a path on at most $\log(4n)$ vertices to two distinct vertices $v_{C_1}, v_{C_2} \in I'$. This implies the existence of a path on at most $2 \log(4n)$ vertices between v_{C_1} and v_{C_2} in G' . Since each component C is $(1, 4)$ -coverable and therefore contains a path on at most eight vertices between any two of its vertices in G , we obtain a path on at most $16 \log(4n)$ vertices between $u \in C_1 \cap I$ and $u' \in C_2 \cap I$, contradicting the independence of I and completing the proof of the theorem. $\square$

11. PROOF OF THEOREM 1.2

For the proof of Theorem 1.2, we need the following theorem of Chudnovsky, Codsi, Fischer and Lokshtanov.

Proposition 11.1 (Theorem 1.2 [CCFL26]). *Let t be a positive integer and let G be a graph that excludes $K_{t,t}$ and $\boxplus_t$ as induced minors and K_t as a subgraph. Then there exists a positive integer $c'(t)$ and $\epsilon'(t) \in (0, 1]$ such that G has a tree decomposition of width at most $2^{c'(t)} \log^{1-\epsilon'(t)} n$.*

Observe that Theorem 1.2 follows immediately from Proposition 11.1 and Theorem 9.2.

Theorem 1.2. *Let t be a positive integer and let G be a graph that excludes $K_{t,t}$ and $\boxplus_t$ as induced minors. Then there exists a positive integer $c(t)$ and $\epsilon(t) \in (0, 1]$, such that G has a tree decomposition (T, χ) such that for every $x \in V(T)$, the distance 8-independence number of $\chi(x)$ is at most $2^{c(t)} \log^{1-\epsilon(t)} n$.*

Proof. Consider a four-radius vertex partition $(X_1, \dots, X_{\pi_4(G)})$ of G with the minimum possible number of parts. We define G' to be the induced minor of G obtained by contracting, for each $i \in [\pi_4(G)]$ and every connected component C of $G[X_i]$, the edges of $G[C]$. In particular, the vertex set of G' is $V(G') = \{v_C \mid C \text{ is a connected component of } G[X_i] \text{ for some } i \in [\pi_4(G)]\}$, and two vertices v_{C_1} and v_{C_2} are adjacent in G' whenever there exist vertices $u_1 \in C_1$ and $u_2 \in C_2$ with $(u_1, u_2) \in E(G)$.

Observe that for every $i \in [\pi_4(G)]$, the set $\{v_C \mid C \text{ is a connected component of } G[X_i]\}$ is an independent set in G' . Hence by Theorem 9.2, the chromatic number of G' , and consequently the size of its largest clique, is at most $\pi_4(G) \leq t \cdot 2^t$. Therefore, by Proposition 11.1, it follows that G' has a tree decomposition (T', χ') such that for every $x \in V(T')$, $|\chi'(x)| \leq 2^{c(t)} \log^{1-\epsilon(t)} n$, where $c(t) := c'(t \cdot 2^t)$ and $\epsilon(t) := \epsilon'(t \cdot 2^t)$.

We define a pair (T, χ) as follows. Let $T := T'$, and for every node $x \in V(T)$ define $\chi(x) := \{u \in V(G) \mid u \in C \subseteq X_i \text{ and } v_C \in \chi'(x)\}$. A direct application of definitions suffices to verify that (T, χ) is indeed a tree decomposition of G . Now assume that there exists a node $x \in V(T)$ such that the distance-8 independence number of $\chi(x)$ is greater than $2^{c(t)} \log^{1-\epsilon(t)} n$. Let $I \subseteq \chi(x)$ be a distance-8 independent set of maximum size. By the pigeonhole principle, there exist distinct vertices $u, v \in I$ and a vertex $v_C \in \chi'(x)$ such that $u, v \in C$, where C denotes the component corresponding to v_C . Since C is $(1, 4)$ -coverable, there exists a path between u and v on at most eight vertices. This contradicts the assumption that I is distance 8-independent, and therefore completes the proof of the theorem. $\square$

12. PROOF OF THEOREM 1.3

For the proof of Theorem 1.3 we require the following three propositions: the first follows directly from a theorem of Bonnet, Hodor, Korhonen and Masařík, the second is the classical theorem of Menger, and the third is due to Gartland, Korhonen, and Lokshtanov.

Proposition 12.1 (Theorem 11 [BHKM23]). *For every positive integer t , there exists a positive integer $\mu(t)$ such that any $K_{t,t}$ -induced-minor-free graph G has a partition $\Lambda := (\Lambda_1, \dots, \Lambda_{\mu(t)})$ of $E(G)$ where, for every $i \in [\mu(t)]$, every connected component of the graph $(V(G), \Lambda_i)$ induces a star in G .*

Proposition 12.2 (Menger's Theorem [Men27]). *Let G be a graph and $A, B \subseteq V(G)$. Then for every positive integer k , G either has a collection of k vertex-disjoint A – B paths, or an A – B separator of size at most k .*

Proposition 12.3 (Theorem 1 [GKL23]). *Let G be a graph of maximum degree Δ and $A, B \subseteq V(G)$. Then for every positive integer k , there either exists a collection of k vertex-disjoint and pairwise anticomplete A – B paths, or there exists an A – B separator of size at most $k \cdot (\Delta + 1)^{\Delta^2+1}$.*

Given a graph G and a positive integer s , an s -size edge partition is defined to be a partition $\Lambda := (\Lambda_1, \dots, \Lambda_\ell)$ of the edge set of G such that every connected component of the graph $(V(G), \Lambda_i)$ has at most s vertices. Proposition 12.1 directly yields the following.

Proposition 12.4. *For every positive integer t , there exists a positive integer $\mu(t)$ such that any $K_{t,t}$ -induced-minor-free graph G of maximum degree Δ has a $(\Delta + 1)$ -size edge partition Λ with at most $\mu(t)$ parts.*

Let G be a graph and $\Lambda := (\Lambda_1, \dots, \Lambda_\ell)$ be a partition of the edge set of G . For $i \in [\ell]$, define the part Λ_i to be *clean* if maximum degree of the graph $(V(G), \Lambda_i)$ is two, and *cluttered* otherwise.

Lemma 12.5. *Let G be a graph and Λ be an s -size edge partition with ℓ parts, of which at least one is cluttered. Let $A, B \subseteq V(G)$ and let f be a positive integer such that G has no A – B separator of size at most f . Then there exists an induced subgraph H of G with an s -size edge partition Λ^H with ℓ parts and strictly more clean parts than Λ , and no A_H – B_H separator of size at most f/s , where $A_H := A \cap V(H)$ and $B_H := B \cap V(H)$.*

Proof. Let Λ_{i_0} be a part in Λ which is cluttered. Let G' be the minor of G obtained by contracting the edges in Λ_{i_0} . Equivalently, $V(G') = \{v_C \mid C \text{ is a connected component of } (V(G), \Lambda_{i_0})\}$, and two vertices v_{C_1} and v_{C_2} are adjacent in G' whenever there exist vertices $u_1 \in C_1$ and $u_2 \in C_2$ with $(u_1, u_2) \in E(G)$. Define $A' := \{v_C \in V(G') \mid C \cap A \neq \emptyset\}$ and $B' := \{v_C \in V(G') \mid C \cap B \neq \emptyset\}$.

Claim 12.5.1. *G' has no A' – B' separator of size at most f/s .*

Proof of the Claim. Suppose otherwise that S' is such a separator, and define $S := \{v \in V(G) \mid v \in C \text{ for some } v_C \in S'\}$. As Λ is an s -size edge partition, each such component C has size at most s , and consequently $|S| \leq s \cdot |S'| \leq f$. Therefore S cannot be an A – B separator, by the definition of G , A and B , which implies the existence of an A – B path P in $G - S$. We obtain a walk P' in $G' - S'$ by replacing each vertex $v \in V(P)$ with the vertex v_C , where C is the connected component of $(V(G), \Lambda_{i_0})$ containing v . By construction of G' , consecutive vertices of P map to vertices in G' that are the same or adjacent, and since P avoids S , the walk P' avoids S' . Moreover, as P starts in A and ends in B , P' starts in A' and ends in B' . But this contradicts the assumption that S' is an A' – B' separator, thereby completing the proof of the claim. $\triangleleft$

Therefore, by Menger's theorem 12.2, there exists a collection $\mathcal{P}'$ of vertex-disjoint A' – B' paths in G' with $|\mathcal{P}'| \geq f/s$. We construct from $\mathcal{P}'$ a collection $\mathcal{P}$ of A – B paths in G as follows. For every path $P' \in \mathcal{P}'$, let $V_{P'} := \{v \in V(G) \mid v \in C \text{ for some } v_C \in V(P')\}$. Since the first vertex of P' lies in

A' and the last lies in B' , the set $V_{P'}$ intersects both A and B . Hence there exists an induced A – B path P in $G[V_{P'}]$, which we add to $\mathcal{P}$. Finally, since the paths in $\mathcal{P}'$ are vertex-disjoint, the sets $V_{P'}$ are pairwise disjoint, and thus the corresponding paths in $\mathcal{P}$ are vertex-disjoint. Consequently, $\mathcal{P}$ is a collection of vertex-disjoint A – B paths in G with $|\mathcal{P}| \geq f/s$.

Let $V_{\mathcal{P}} := \{v \in V(G) \mid v \in P \text{ for some } P \in \mathcal{P}\}$ and let $H := G[V_{\mathcal{P}}]$. For each $i \in [\ell]$, define $\Lambda_i^H := \{(u, v) \in \Lambda_i \mid u, v \in V_{\mathcal{P}}\}$. We claim that $\Lambda^H := (\Lambda_1^H, \dots, \Lambda_\ell^H)$ satisfies the requirements of the lemma. By definition, Λ^H is a partition of $E(H)$ into ℓ parts. Since $\Lambda_i^H \subseteq \Lambda_i$ and $V(H) \subseteq V(G)$, every connected component of $(V(H), \Lambda_i^H)$ is contained in a component of $(V(G), \Lambda_i)$. Hence Λ^H is an s -size edge partition.

Now fix $i \in [\ell]$ such that Λ_i is clean. If a vertex $v \in V(H)$ has at least three neighbors in $(V(H), \Lambda_i^H)$, then it has the same neighbors in $(V(G), \Lambda_i)$, contradicting the definition of i . Hence Λ_i^H is clean. Finally, consider the graph $(V(H), \Lambda_{i_o}^H)$ and suppose it contains a vertex v of degree greater than two. Let P be the path in $\mathcal{P}$ containing v . Since P is induced, at most two neighbors of v are in P and hence at least one neighbor of v in $(V(H), \Lambda_{i_o}^H)$ must belong to a different path Q . Thus both P and Q intersect the connected component C of $(V(G), \Lambda_{i_o})$ that contains v . But then, P' and Q' must both contain v_C , which contradicts the fact that they were vertex-disjoint. Therefore $\Lambda_{i_o}^H$ is clean, completing the proof. $\square$

Lemma 12.6. *Let G be a graph with an s -size edge partition Λ with ℓ parts, of which z are cluttered, for positive integers s, z, ℓ . Let k be a positive integer and let $A, B \subseteq V(G)$. Then G either has a collection of k vertex-disjoint and pairwise anticomplete A – B paths, or has an A – B separator of size at most $k \cdot g(s, z, \ell)$, where $g(s, z, \ell) := s^z \cdot (2\ell + 1)^{4\ell^2 + 1}$.*

Proof. We prove the lemma by induction on z . For the base case, observe that when z is zero, every part in Λ is clean and consequently the maximum degree of G is at most 2ℓ . Applying Proposition 12.3 to G asserts the lemma for $z = 0$. Fix $z_o \geq 1$ and assume that the lemma holds for all instances with $z < z_o$. If G has an A – B separator of size at most $k \cdot g(s, z_o, \ell)$, then the lemma holds and we are done. Otherwise, we apply Lemma 12.5 to (G, Λ, A, B) and obtain an induced subgraph H together with an s -size edge partition Λ^H with ℓ parts, of which fewer than z_o are cluttered. Let $A_H := A \cap V(H)$ and $B_H := B \cap V(H)$, and apply the inductive hypothesis to the instance $(H, \Lambda^H, k, A_H, B_H)$. Since H admits no A_H – B_H separator of size at most $\frac{k \cdot g(s, z_o, \ell)}{s} \leq k \cdot g(s, z_o - 1, \ell)$, we conclude that H contains a collection of k vertex-disjoint and pairwise anticomplete paths. As H is an induced subgraph of G , these paths remain vertex-disjoint and pairwise anticomplete in G , which concludes the proof. $\square$

Proposition 12.4, combined with the fact that the number of cluttered parts is at most the number of parts, immediately yields the following corollary.

Corollary 12.7. *Let Δ, t be positive integers and let G be a $K_{t,t}$ -induced-minor-free graph of maximum degree at most Δ . Let k be a positive integer and let $A, B \subseteq V(G)$. Then G either has a collection of k vertex-disjoint and pairwise anticomplete A – B paths, or an A – B separator of size at most $k \cdot g(\Delta, t)$, where $g(\Delta, t) := (\Delta + 1)^\mu \cdot (2\mu + 1)^{4\mu^2 + 1}$ and μ is the constant defined in Proposition 12.4.*

We now proceed to the proof of Theorem 1.3, which we restate for convenience.

Theorem 1.3. *Let t be a positive integer and let G be a $K_{t,t}$ -induced-minor-free graph. Let $A, B \subseteq V(G)$ and k be a positive integer. Then G either has a collection of k vertex-disjoint and pairwise anticomplete A – B paths or an A – B separator $S \subseteq V(G)$ whose distance $16 \log(4n)$ -independence number is $100k \cdot g \cdot \log^3(4n)$, where $g := (24 \log^2(4n) + 4d + 1)^\mu \cdot (2\mu + 1)^{4\mu^2 + 1}$ and d, μ are the constants defined in Proposition 10.3 and 12.4, respectively.*

Proof. Consider a four-radius vertex partition $(X_1, \dots, X_{\pi_4(G)})$ of G with at most $t \cdot 2^t$ parts, whose existence is asserted by Theorem 9.2. We define G' to be the induced minor of G obtained by contracting, for each $i \in [\pi_4(G)]$, all edges inside the connected components of $G[X_i]$. Explicitly, the vertex set of G' is $V(G') = \{v_C \mid C \text{ is a connected component of } G[X_i], \text{ for some } i \in [\pi_4(G)]\}$, and two vertices v_{C_1} and v_{C_2} are adjacent in G' whenever there exist vertices $u_1 \in C_1$ and $u_2 \in C_2$ with $(u_1, u_2) \in E(G)$. We use n' to denote the number of vertices in G' .

Since G' is obtained from a $K_{t,t}$ -induced-minor-free graph by contracting edges, it is itself $K_{t,t}$ -induced-minor-free. Also, the set $\{v_C \mid C \text{ is a connected component of } G[X_i]\}$ is an independent set in G' , for every $i \in [\pi_4(G)]$. So it follows that the chromatic number, and hence the clique number, of G' is at most $\pi_4(G) \leq t \cdot 2^t$. Hence, by Proposition 10.3, G' has degeneracy at most $d \leq d(t \cdot 2^t)$. Define $A' := \{v_C \in V(G') \mid C \cap A \neq \emptyset\}$ and $B' := \{v_C \in V(G') \mid C \cap B \neq \emptyset\}$. Let $\mathcal{F}$ be the $4d$ -neighborhood-witnessing layered family of G' with thickness at most $\log(4n')$, whose existence is guaranteed by Lemma 8.1. Applying Theorem 8.7 to $(G', A', B', \mathcal{F})$ with $f := 12k \cdot g + 1$ yields either a vertex set S' or an induced subgraph H' of G' .

Suppose we obtain an induced subgraph H' of maximum degree at most $24 \log^2(4n') + 4d$ such that every $A_{H'}-B_{H'}$ separator S in H' satisfies $\text{cov}_{\mathcal{F}}(S) > k \cdot g$, where $A_{H'} := A' \cap V(H')$ and $B_{H'} := B' \cap V(H')$. Applying Corollary 12.7 to $(H', A_{H'}, B_{H'}, k)$ yields either a collection of k vertex-disjoint and pairwise anticomplete $A_{H'}-B_{H'}$ paths, or an $A_{H'}-B_{H'}$ separator of size at most $k \cdot g$. The latter contradicts the definition of H' , and therefore H' contains a collection $\mathcal{P}'$ of k vertex-disjoint and pairwise anticomplete $A_{H'}-B_{H'}$ paths. We now lift $\mathcal{P}'$ to a collection $\mathcal{P}$ of $A-B$ paths in G . For each path $P' \in \mathcal{P}'$, let $V_{P'} := \{v \in V(G) \mid v \in C \text{ for some } v_C \in V(P')\}$. Since P' starts in $A_{H'}$ and ends in $B_{H'}$, the set $V_{P'}$ intersects both A and B , and hence contains an induced $A-B$ path P , which we add to $\mathcal{P}$. As the paths in $\mathcal{P}'$ are vertex-disjoint and pairwise anticomplete, the sets $V_{P'}$ are likewise disjoint and anticomplete, and so the corresponding paths in $\mathcal{P}$ also have these properties. Consequently, $\mathcal{P}$ is a collection of k vertex-disjoint and pairwise anticomplete $A-B$ paths in G .

Otherwise we obtain an $A'-B'$ separator $S' \subseteq V(G')$ in G' , such that $\text{fcov}_{\mathcal{F}}(S') \leq 8 \log^2(4n') \cdot (12k \cdot g + 1)$. By Proposition 8.3, we conclude that $\text{cov}_{\mathcal{F}}(S') \leq 100k \cdot g \cdot \log^3(4n')$. Let $\mathcal{F}' \subset \mathcal{F}$ be a cover of S' of size at most $100k \cdot g \cdot \log^3(4n')$, and redefine $S' := \bigcup_{F \in \mathcal{F}'} F$. Let $S := \{v \in V(G) \mid v \in C, \text{ for some } v_C \in S'\}$. We claim that S is an $A-B$ separator with the desired properties. First, observe that if it is not an $A-B$ separator, then $G - S$ must contain an $A-B$ path P . We obtain a walk P' in $G' - S'$ by replacing each vertex $v \in V(P)$ with the vertex v_C , where for some $i \in [\pi_4(G)]$, C is the connected component of $G[X_i]$ containing v . By construction of G' , consecutive vertices of P map to vertices in G' that are the same or adjacent, and since P avoids S , the walk P' avoids S' . Moreover, as P starts in A and ends in B , P' starts in A' and ends in B' . But this contradicts the assumption that S' is an $A'-B'$ separator in G' . Second, observe that since every connected component C of $G[X_i]$ for each $i \in [\pi_4(G)]$ is $(1, 4)$ -coverable in G , for every pair of vertices $u, v \in C$ there exists a $u-v$ path on at most eight vertices in G . Furthermore, letting $\mathcal{C}_{\mathcal{F}'}$ denote the set of centers of the elements of $\mathcal{F}'$, every vertex $v_C \in S'$ has a path on at most $\log(4n')$ vertices in G' to some vertex of $\mathcal{C}_{\mathcal{F}'}$. Let $\mathcal{C} \subseteq V(G)$ be a subset that contains one arbitrary vertex of C for every $v_C \in \mathcal{C}_{\mathcal{F}'}$. Combining the above two observations, we conclude that every vertex of S has a path on at most $8 \log(4n') \leq 8 \log(4n)$ vertices to some vertex of $\mathcal{C}$. Therefore, the maximum size of a distance $16 \log(4n)$ -independent set in S is $|C| \leq 100k \cdot g \cdot \log^3(4n)$, which completes the proof of the theorem. $\square$

REFERENCES

- [ABVT25] Bogdan Alecu, Édouard Bonnet, Pedro Bureo Villafana, and Nicolas Trotignon. Every graph is essential to large treewidth. *CoRR*, abs/2502.14775, 2025.

- [ACHS22] Tara Abrishami, Maria Chudnovsky, Sepehr Hajebi, and Sophie Spirkl. Induced subgraphs and tree decompositions III. Three-path-configurations and logarithmic treewidth. *Advances in Combinatorics*, sep 9 2022.
- [ACK⁺25] Tara Abrishami, Jadwiga Czyżewska, Kacper Kluk, Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzażewski. On coarse tree decompositions and coarse balanced separators, 2025.
- [AHJ⁺24] Sandra Albrechtsen, Tony Huynh, Raphael W. Jacobs, Paul Knappe, and Paul Wollan. A menger-type theorem for two induced paths. *SIAM Journal on Discrete Mathematics*, 38(2):1438–1450, May 2024.
- [BBCD24] Romain Bourneuf, Matija Bucić, Linda Cook, and James Davies. On polynomial degree-boundedness. *Advances in Combinatorics*, oct 9 2024.
- [BHKM23] Édouard Bonnet, Jędrzej Hodor, Tuukka Korhonen, and Tomáš Masařík. Treewidth is polynomial in maximum degree on weakly sparse graphs excluding a planar induced minor. *arXiv preprint arXiv:2312.07962*, 2023.
- [Bod98] Hans L Bodlaender. A partial k-arboretum of graphs with bounded treewidth. *Theoretical computer science*, 209(1-2):1–45, 1998.
- [CC16] Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem. *Journal of the ACM (JACM)*, 63(5):1–65, 2016.
- [CCFL26] Maria Chudnovsky, Julien Codsí, David Fischer, and Daniel Lokshtanov. Induced minors and subpolynomial treewidth, 2026.
- [CCL⁺25] Maria Chudnovsky, Julien Codsí, Daniel Lokshtanov, Martin Milanič, and Varun Sivashankar. Tree independence number v. walls and claws, 2025.
- [CEL25] Maria Chudnovsky, Ajaykrishnan E S, and Daniel Lokshtanov. (treewidth, clique)-boundedness and polylogarithmic tree-independence, 2025.
- [CFK⁺15] Marek Cygan, Fedor V Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized algorithms*, volume 5. Springer, 2015.
- [CGH⁺25] Maria Chudnovsky, Peter Gartland, Sepehr Hajebi, Daniel Lokshtanov, and Sophie Spirkl. Tree independence number IV. Even-hole-free graphs. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4444–4461. SIAM, Philadelphia, PA, 2025.
- [Chv79] Vasek Chvatal. A greedy heuristic for the set-covering problem. *Mathematics of operations research*, 4(3):233–235, 1979.
- [CT21] Julia Chuzhoy and Zihan Tan. Towards tight (er) bounds for the excluded grid theorem. *Journal of Combinatorial Theory, Series B*, 146:219–265, 2021.
- [Die25] Reinhard Diestel. *Graph theory*. Springer Nature, 2025.
- [DMŠ24] Clément Dallard, Martin Milanič, and Kenny Štorgel. Treewidth versus clique number. ii. tree-independence number. *Journal of Combinatorial Theory, Series B*, 164:404–442, 2024.
- [GH25] António Girão and Zach Hunter. Induced subdivisions in ks,s-free graphs with polynomial average degree. *International Mathematics Research Notices*, 2025(4):rnaf025, 02 2025.
- [GKL23] Peter Gartland, Tuukka Korhonen, and Daniel Lokshtanov. On induced versions of menger’s theorem on sparse graphs, 2023.
- [GP25] Agelos Georgakopoulos and Panos Papasoglu. Graph minors and metric spaces. *Combinatorica*, 45(3), June 2025.
- [Haj25] Sepehr Hajebi. Polynomial bounds for pathwidth, 2025.
- [HNST24] Kevin Hendrey, Sergey Norin, Raphael Steiner, and Jérémie Turcotte. On an induced version of menger’s theorem. *The Electronic Journal of Combinatorics*, 31(4), November 2024.
- [HR90] Torben Hagerup and Christine Rüb. A guided tour of chernoff bounds. *Information processing letters*, 33(6):305–308, 1990.
- [KLS⁺24] Viktoriia Korchemna, Daniel Lokshtanov, Saket Saurabh, Vaishali Surianarayanan, and Jie Xue. Efficient approximation of fractional hypertree width. In *65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024*, pages 754–779. IEEE, 2024.
- [LMM⁺24] Paloma T. Lima, Martin Milanec, Peter Mursic, Karolina Okrasa, Paweł Rzażewski, and Kenny Štorgel. Tree decompositions meet induced matchings: Beyond max weight independent set. In *32nd Annual European Symposium on Algorithms, ESA 2024, Royal Holloway, London, United Kingdom, September 2-4, 2024*, volume 308 of *LIPICs*, pages 85:1–85:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.
- [LPR26] Daniel Lokshtanov, Michał Pilipczuk, and Paweł Rzażewski. Finding large sparse induced subgraphs in graphs of small (but not very small) tree-independence number, 2026.
- [LR99] Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM (JACM)*, 46(6):787–832, 1999.
- [Men27] Karl Menger. Zur allgemeinen Kurventheorie. *Comptes Rendus de la Société des Sciences et des Lettres de Varsovie*, 20:1–3, 1927.

- [MU17] Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press, 2017.
- [NSS25a] Tung Nguyen, Alex Scott, and Paul Seymour. Asymptotic structure. i. coarse tree-width, 2025.
- [NSS25b] Tung Nguyen, Alex Scott, and Paul Seymour. A counterexample to the coarse menger conjecture, 2025.
- [RS86] Neil Robertson and Paul Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986.
- [ST21] Ni Luh Dewi Sintiari and Nicolas Trotignon. (theta, triangle)-free and (even hole, $k-4$)-free graphs—part 1: Layered wheels. *Journal of Graph Theory*, 97(4):475–509, 2021.
- [Won83] Richard T. Wong. Combinatorial optimization: Algorithms and complexity (christos h. papadimitriou and kenneth steiglitz). *SIAM Review*, 25(3):424–425, 1983.