

Coarse Balanced Separators in Fat-Minor-Free Graphs

Édouard Bonnet*

CNRS, ENS de Lyon,
Université Claude Bernard Lyon 1,
LIP, UMR 5668
edouard.bonnet@ens-lyon.fr

Hung Le†

University of Massachusetts Amherst
hungle@cs.umass.edu

Marcin Pilipczuk‡

University of Warsaw
marcin.pilipczuk@mimuw.edu.pl

Michał Pilipczuk§

University of Warsaw
michal.pilipczuk@mimuw.edu.pl

Abstract

Fat minors are a coarse analogue of graph minors where the subgraphs modeling vertices and edges of the embedded graph are required to be distant from each other, instead of just being disjoint. In this paper, we give a coarse analogue of the classic theorem that an n -vertex graph excluding a fixed minor admits a balanced separator of size $\mathcal{O}(\sqrt{n})$. Specifically, we prove that for every integer d , real $\varepsilon > 0$, and graph H , there exist constants c and r such that every n -vertex graph G excluding H as a d -fat minor admits a set $S \subseteq V(G)$ that is a balanced separator of G and can be covered by $cn^{\frac{1}{2}+\varepsilon}$ balls of radius r in G . Our proof also works in the weighted setting where the balance of the separator is measured with respect to any weight function on the vertices, and is effective: we obtain a randomized polynomial-time algorithm to compute either such a balanced separator, or a d -fat model of H in G .

1 Introduction

Coarse graph theory is a relatively young direction in structural graph theory, whose aim is to describe the structure in graphs understood as metric spaces. As suggested in one of the foundational works of Georgakopoulos and Papasoglu [10], the hope is that many classic notions and results of structural graph theory, particularly the theory of graph minors, could be lifted to suitable coarse analogues. The usual way of obtaining such a coarse analogue is to replace any constraints of disjointness or intersection with the requirement that the relevant objects are far or close to each other, respectively. A good example of such a lift is the coarse analogue of the notion of a minor, called a *fat minor*. Formally, for graphs G and H , we say that a graph G contains H as a d -fat minor, for a parameter $d \in \mathbb{N}$, if there are connected subgraphs $\{B_u : u \in V(H)\}$ and $\{B_e : e \in E(H)\}$ of G such that

- whenever u is an endpoint of e , the subgraphs B_u and B_e intersect; and
- except for the above, all the subgraphs in $\{B_f : f \in V(H) \cup E(H)\}$ are at distance at least d from each other in G .

*Édouard Bonnet was supported by the French National Research Agency through the project TWIN-WIDTH with reference number ANR-21-CE48-0014.

†Hung Le was supported by the NSF CAREER award CCF-2237288 and NSF grant CCF-2517033.

‡The work of Marcin Pilipczuk on this manuscript was supported by the project BUKA that has received funding from the European Research Council (ERC), grant agreement No. 101126229.

§The work of Michał Pilipczuk on this manuscript was supported by the project BOBR that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 948057).



European Research Council
Established by the European Commission



The collection of subgraphs $\{B_u : u \in V(H)\}$ and $\{B_e : e \in E(H)\}$ as above is called a *d-fat model* of H in G . Thus, the notion of a fat minor requires the constituent parts of the model to be far from each other, rather than just disjoint as in the classic notion of a minor.

In the light of all the deep advances of the classic theory of graph minors, the following question becomes natural: Supposing G excludes a fixed graph H as a d -fat minor, what can be said about the metric structure of G ? Georgakopoulos and Papasoglu [10] conjectured that such graphs are *quasi-isometric* to graphs excluding H as a classic minor. Roughly speaking, this would mean that the metric space of a graph excluding H as a d -fat minor can be mapped to a metric space of some H -minor-free graph G' so that the mapping preserves the distances up to a constant multiplicative and additive factor (both depending on d and H). If this were true, then the wealth of known properties of H -minor-free graphs could be pulled through the quasi-isometry to their coarse analogues that would apply to graphs excluding H as a d -fat minor.

Unfortunately, as shown by Davies, Hickingbotham, Illingworth, and McCarty [7], the conjecture of Georgakopoulos and Papasoglu turns out to be false; see also the newer works [2, 3] for simpler and stronger counterexamples. However, many of its potential consequences — structural properties of graphs excluding a fat minor — can still be true, at least to some extent. In this work we identify one such property: the existence of balanced separators of sublinear size.

To measure the balance of a separator, it will be convenient to speak about vertex-weighted graphs. A *weighted graph* is a graph G equipped with a vertex weight function $\mathbf{w}_G : V(G) \rightarrow \mathbb{R}_{\geq 0}$. (We often omit the subscript if G is clear from the context.) A *balanced separator* of G is a set $S \subseteq V(G)$ such that for every connected component C of $G - S$, we have $\mathbf{w}(C) \leq \mathbf{w}(G)/2$. Here, if X is a subgraph or a vertex subset of G , then $\mathbf{w}(X)$ denotes the total weight of the vertices featured in X .

It is well-known that H -minor-free graphs admit balanced separators of square root size.

Theorem 1 ([4]). *For every graph H , every n -vertex weighted graph that excludes H as a minor has a balanced separator of size $\mathcal{O}_H(\sqrt{n})$.¹*

Note that for planar graphs (which coincide with $\{K_5, K_{3,3}\}$ -minor-free graphs), this is exactly the Planar Separator Theorem of Lipton and Tarjan [15]. In general, the existence of square-root-sized balanced separators is a fundamental reason behind countless combinatorial and algorithmic results applicable to planar and H -minor-free graphs.

A natural coarse analogue of **Theorem 1** would be the following statement, which we pose as a conjecture. Here, for a graph G and a vertex subset $S \subseteq V(G)$, we say that S is *(k, r)-coverable* if S can be covered by k balls of radius r in G ; or in other words, there exists $X \subseteq V(G)$ with $|X| \leq k$ such that every vertex of S is at distance at most r from some vertex of X .

Conjecture 1. *For every graph H and $d \in \mathbb{N}$, there exist constants $c, r \in \mathbb{N}$ such that every n -vertex weighted graph G that excludes H as a d -fat minor has a $(c\sqrt{n}, r)$ -coverable balanced separator.*

Our results. Our main result is a weaker form of **Conjecture 1**, where we lose a multiplicative factor n^ε in the number of balls covering the separator, for any fixed $\varepsilon > 0$. In the formulation below we make the dependence on H , d , and ε explicit, as it is actually not too terrible.

Theorem 2. *For every graph H , integer $d \in \mathbb{N}$, and real $\varepsilon > 0$, the following holds: Every n -vertex weighted graph that excludes H as a d -fat minor has a balanced separator that is $(\mathcal{O}(\|H\|^2 \cdot n^{1/2+\varepsilon}), \mathcal{O}(d/\varepsilon))$ -coverable, where $\|H\| = |V(H)| + |E(H)|$.*

Furthermore, there is a randomized polynomial-time algorithm that, given a weighted graph G , either returns such a balanced separator in G or finds a d -fat minor model of H in G .

We remark that while several previous works studied the structure of H -fat-minor-free graphs for specific graphs H , to the best of our knowledge, **Theorem 2** offers the first non-trivial structural property of H -fat-minor-free graphs proved for an *arbitrary* H . The algorithm of **Theorem 2** can actually report

¹For a tuple of parameters \bar{p} , the $\mathcal{O}_{\bar{p}}(\cdot)$ notation hides factors that may depend on \bar{p} .

the centers of the balls covering the separator (rather than the separator itself); alternatively, the centers can be obtained greedily from the separator at the unsequential expense of doubling the radius. We also note that by an observation of Abrishami et al. [1], **Theorem 2** implies that an n -vertex graph excluding H as a d -fat minor has *coarse treewidth* $(\mathcal{O}(\|H\|^2 \cdot n^{1/2+\varepsilon}), \mathcal{O}(d/\varepsilon))$, for any fixed $\varepsilon > 0$; we refer to [1] for relevant definitions.

As a side result, we also prove a variant for excluding induced minors in which the balanced separator can be covered by balls of radius 1. Such balanced separators are often called *dominated balanced separators*, and they seem to play an important role in the emerging structure theory for induced minors. Recall here that H is an *induced minor* of G if there are pairwise disjoint connected subgraphs $\{B_u\}_{u \in V(H)}$ in G such that for any pair of distinct vertices $u, v \in V(H)$, we have that u and v are adjacent in H if and only if B_u and B_v are adjacent in G .

Theorem 3. *For every graph H , every n -vertex weighted graph that is H -induced-minor-free has an $(\mathcal{O}_H(n^{2/3} \log^{4/3} n), 1)$ -coverable balanced separator.*

We anticipate that the bound provided by **Theorem 3** is not tight and the constant $2/3$ in the exponent can be improved to $1/2$. Indeed, we note that this can be done when H is a complete bipartite graph, at the cost of increasing the radius to 4.

Theorem 4. *For every $t \in \mathbb{N}$, every n -vertex weighted graph that is $K_{t,t}$ -induced-minor-free has an $(\mathcal{O}_t(n^{1/2} \log n), 4)$ -coverable balanced separator.*

Our techniques. Our proof of **Theorem 2** is inspired by the following result of Korhonen and Lokshtanov [11] about the existence of balanced separators in induced-minor-free graphs.

Theorem 5 ([11, Theorem 1.1]). *For every graph H , every m -edge weighted graph G that is H -induced-minor-free has a balanced separator of size at most $\mathcal{O}_H(\sqrt{m} \log m)$.*

We note that in [11], this result is stated for unweighted graphs, but the proof can be easily lifted to the weighted setting. (Alternatively, the weighted result can be obtained from the unweighted one using a theorem of Dvořák and Norin [8, Theorem 1].) Also, the formulation in [11] postulates a stronger bound of $\mathcal{O}_H(\sqrt{m})$, but the proof is partly based on a faulty lemma from an earlier work of Lee [13]. Without reliance on this lemma, the proof goes through with an extra $\log m$ factor.

The idea of Korhonen and Lokshtanov in the proof of **Theorem 5** is to exploit the classic result of Leighton and Rao [14] on the gap between product multicommodity flows and balanced separators. Roughly speaking, from this result it follows that if a weighted graph G does not admit a balanced separator of small size, then in G there is a multicommodity flow λ that pushes, for every pair of vertices u, v , a significant amount of flow from u to v , and at the same time every vertex of G carries only a relatively small amount of flow. Such a flow λ can be understood as a well-distributed system of paths in G . Korhonen and Lokshtanov show that by sampling paths from this system, one can find, with positive probability, an induced minor model of the supposedly excluded graph H , leading to a contradiction.

In our proof of **Theorem 2**, we show that assuming the graph is suitably sparse, the same strategy can be used to expose a fat minor model of H , not only an induced one. The precise notion of sparseness has to be chosen carefully. We achieve it using the following clustering result of Filtser [9]: for every graph G and $\varepsilon > 0$, one can partition the vertex set of G into subsets inducing subgraphs of diameter $\mathcal{O}(1/\varepsilon)$, called *clusters*, so that every ball of radius 2 in G intersects only $\mathcal{O}(n^\varepsilon)$ clusters. The Leighton–Rao argument is then applied to the graph G' obtained from G by contracting every cluster to a single vertex. Thus, a balanced separator of G' of size $\mathcal{O}_H(n^{1/2+\varepsilon})$ can be naturally lifted to a balanced separator of G that can be covered by $\mathcal{O}_H(n^{1/2+\varepsilon})$ balls of bounded radius.

The proof of **Theorem 3** is much less involved. We first apply a simple clustering scheme that vertex-partitions the graph into connected subgraphs, each dominated by a single vertex, so that contracting the subgraphs yields a graph with $\mathcal{O}_H(n^{4/3} \log^{2/3} n)$ edges. Then we apply **Theorem 5**

as a black-box in the contracted graph, and again lift the obtained balanced separator to the original graph. **Theorem 4** follows by the same approach, but we apply an off-the-shelf clustering result for $K_{t,t}$ -induced-minor-free graphs due to Chudnovsky, Codsi, Ajaykrishnan, and Lokshtanov [6].

2 Preliminaries

For a nonnegative integer k , we denote $[k] := \{1, \dots, k\}$.

Graphs. All graphs considered in this paper are finite, simple (with no parallel edges or loops), and undirected, unless explicitly stated. We use standard graph notation. In particular, for a graph G , we write $V(G)$ for the vertex set of G , $E(G)$ for the edge set of G , $\text{dist}_G(\cdot, \cdot)$ for the distance metric in G (mapping every pair of vertices u, v to the length, i.e., number of edges, on a shortest path between u and v), and $\text{diam}(G)$ for the *diameter* of G , which is the largest distance between a pair of vertices of G . When X, Y are subgraphs or subsets of vertices of G , we define $\text{dist}_G(X, Y)$ as the smallest distance in G between a vertex of X and a vertex of Y . For a graph G and a vertex or edge x , we write $x \in G$ to signify that x belongs to the vertex set or the edge set of G , respectively. For a pair of graphs G, G' , we write $G \cup G'$ for the graph $(V(G) \cup V(G'), E(G) \cup E(G'))$.

For $r \in \mathbb{N}$ and a graph G , the *ball* of radius r with center $u \in V(G)$ is the set $B_G(u, r) := \{v \in V(G) \mid \text{dist}_G(u, v) \leq r\}$. The *r th power* of G is the graph G^r on the same vertex set as G where two vertices u, v are adjacent if and only if they are at distance at most r in G .

For a (possibly directed) graph G and a pair of distinct vertices $u, v \in V(G)$, we let $\mathcal{P}_{u,v}(G)$ be the set of all paths from u to v in G . Then $\mathcal{P}(G) := \bigcup_{u,v \in V(G), u \neq v} \mathcal{P}_{u,v}(G)$ is the set of all paths in G on at least two vertices. If G is directed, then we mean the paths as directed.

Weighted graphs and balanced separators have already been defined in [Section 1](#). Note that we use vertex weights only to measure the balance of separators. In particular, there are no weights on the edges and the lengths of paths are measured by the number of edges. When taking (induced) subgraphs in the context of weighted graphs, we naturally restrict the weight function to the remaining vertices.

A *connected partition* of a graph G is a partition \mathcal{X} of the vertex set of G such that for each set $X \in \mathcal{X}$, the graph $G[X]$ is connected. Sets $X \in \mathcal{X}$ are called the *clusters* of \mathcal{X} . The *strong diameter* of \mathcal{X} is $\max_{X \in \mathcal{X}} \text{diam}(G[X])$, and the *weak diameter* of \mathcal{X} is $\max_{X \in \mathcal{X}} \max_{x,y \in X} \text{dist}_G(x, y)$. Finally, the *quotient graph* G/\mathcal{X} is the graph with vertex set \mathcal{X} where two clusters $X, Y \in \mathcal{X}$ are adjacent if in G there is an edge with one endpoint in X and the other in Y . In other words, G/\mathcal{X} is obtained from G by contracting every cluster into a single vertex and removing parallel edges and self-loops.

We will later use the following lemma, a consequence of the work of Filtser.

Lemma 1 (Consequence of [9, Theorem 4]). *There is a randomized polynomial-time algorithm that, given a graph G and a real $\varepsilon > 0$, outputs a connected partition \mathcal{X} of G of strong diameter at most $\frac{32}{\varepsilon}$ such that for every vertex $u \in V(G)$, the ball $B_G(u, 2)$ intersects $\mathcal{O}(n^\varepsilon)$ clusters of \mathcal{X} .*

Proof. In [9, Theorem 4], set $\Delta := \frac{32}{\varepsilon}$ and $\alpha := \frac{2}{\varepsilon}$. □

Minors. The notions of minors, fat minors, and induced minors were also already introduced in [Section 1](#). Similarly to the work of Korhonen and Lokshtanov [11] on induced minors, it will be useful to work with the following relaxed notion of fat models. A *crude d -fat model* of a graph H in a graph G is a pair of mappings (ϕ, π) such that ϕ maps vertices of H to vertices of G , π maps edges of H to paths in G , and the following conditions hold:

- For every edge $uv \in E(H)$, $\pi(uv)$ is a $\phi(u)$ – $\phi(v)$ path in G .
- For every pair of edges $uv, u'v' \in E(H)$ with $|\{u, v, u', v'\}| = 4$, we have $\text{dist}_G(\pi(uv), \pi(u'v')) \geq d$.

Note that we do not require the vertices $\phi(u)$, for $u \in V(H)$, to be far from each other, or even distinct. The only distance constraints are between paths modeling pairs of edges that do not share endpoints.

The following simple lemma, directly lifted from an analogous statement in [11], explains that a d -fat model of a graph H can be easily extracted from a crude d -fat model of the **2-subdivision** \check{H} of H — the graph obtained from H by replacing every edge by a path of length 3.

Lemma 2. *Let H be a graph with no isolated vertices. Suppose G contains a crude d -fat model of \check{H} . Then G also contains a d -fat model of H .*

Proof. Let (ϕ, π) be a crude d -fat model of \check{H} in G . For every vertex $u \in V(H)$, let \check{E}_u be the set of all edges incident to u in \check{H} . We set

$$B_u := \bigcup_{e \in \check{E}_u} \pi(e).$$

For every edge $e \in E(H)$, let $\check{e} \in E(\check{H})$ be the middle edge of the 3-edge-path into which e got subdivided in \check{H} . We set

$$B_e := \pi(\check{e}).$$

We now show that $(\{B_u\}_{u \in V(H)}, \{B_e\}_{e \in E(H)})$ is a d -fat model of H in G . First, note that every subgraph B_u and B_e is indeed connected. Second, observe that for every vertex u of H and edge e of H incident to u , the subgraph B_u and B_e intersect, because they both contain $\phi(u')$, where u' is the endpoint of \check{e} that is adjacent to u in \check{H} .

We are left with verifying the distance condition. Let X, Y be two members of the collection of subgraphs $\{B_u\}_{u \in V(H)} \cup \{B_e\}_{e \in E(H)}$. We have to consider three cases.

Case 1: Suppose $X = B_u$ and $Y = B_v$ for distinct $u, v \in V(H)$. Let $x \in B_u$ and $y \in B_v$ be any two vertices. Then, by construction, we have $x \in \pi(e_x)$ and $y \in \pi(e_y)$ for some edges $e_x, e_y \in E(\check{H})$ incident to u and v , respectively. Observe that e_x and e_y do not share any endpoint, and therefore we have $\text{dist}_G(\pi(e_x), \pi(e_y)) \geq d$. Thus $\text{dist}_G(x, y) \geq d$, and hence also $\text{dist}_G(B_u, B_v) \geq d$.

Case 2: Suppose $X = B_u$ and $Y = B_e$ for some edge e that is not incident to u . Then for every edge $f \in \check{E}_u$, f and \check{e} do not share any endpoint. It follows that $\text{dist}_G(\pi(f), \pi(\check{e})) \geq d$, and hence also $\text{dist}_G(B_u, B_e) \geq d$.

Case 3: Suppose $X = B_e$ and $Y = B_f$ for some distinct edges e and f . Then \check{e} and \check{f} do not share any endpoint and therefore, $\text{dist}_G(B_e, B_f) = \text{dist}_G(\pi(\check{e}), \pi(\check{f})) \geq d$. \square

We note that the construction from the proof of **Lemma 2** can be easily turned into a polynomial-time algorithm for converting a crude d -fat model of \check{H} into a d -fat model of H .

Concurrent flows and balanced separators. Our proof of **Theorem 2** crucially relies on the duality between multicommodity flows and balanced cuts, proved by Leighton and Rao [14]. Let us recall the relevant results.

The work of Leighton and Rao concerns the **Product Multicommodity Flow Problem** (PMFP) on directed edge-capacitated graphs. An instance of (directed) PMFP is a triple (D, c, π) where

- D is a (directed) graph,
- $c: E(D) \rightarrow \mathbb{R}_{\geq 0}$ is an edge-capacity function, and
- $\pi: V(D) \rightarrow \mathbb{R}_{\geq 0}$ is a vertex demand function.

For every pair of distinct vertices $u, v \in V(D)$, the **demand** from u to v is $\pi(u)\pi(v)$ units of commodity (u, v) . The goal is to find a (multicommodity) flow — union of flows of commodity (u, v) for every pair of distinct vertices $u, v \in V(D)$ — that does not exceed the capacity of any edge and routes for all these pairs simultaneously $q \cdot \pi(u)\pi(v)$ units of the commodity (u, v) from u to v , for the largest

possible value q . Formally, a multicommodity flow of value $q \in \mathbb{R}_{\geq 0}$ is a function $f : \mathcal{P}(D) \rightarrow \mathbb{R}_{\geq 0}$, assigning each path $P \in \mathcal{P}(D)$ the amount of flow passing through P , so that

$$\sum_{P \in \mathcal{P}(D), e \in P} f(P) \leq c(e) \quad \text{for all } e \in E(D) \quad \text{and} \quad \sum_{P \in \mathcal{P}_{u,v}(D)} f(P) = q \cdot \pi(u)\pi(v) \quad \text{for all } u, v \in V(D), u \neq v.$$

On the other hand of the duality we have balanced cuts, where the balance is measured as follows: For a partition (U, \bar{U}) of $V(D)$ such that $\pi(U)\pi(\bar{U}) > 0$ (with $\pi(X) := \sum_{v \in X} \pi(v)$), we define its *weighted ratio cost* as

$$\frac{\sum_{e \in E_D(U, \bar{U})} c(e)}{\pi(U) \cdot \pi(\bar{U})},$$

where $E_D(U, \bar{U})$ denotes the set of edges of D with tail in U and head in \bar{U} . With these definitions in place, we can recall the main result of Leighton and Rao.

Theorem 6 ([14, Theorem 17]). *There is a randomized polynomial-time algorithm that given an instance (D, c, π) of the Directed Product Multicommodity Flow Problem and a real $q \in \mathbb{R}_{\geq 0}$, outputs either*

- (1) a multicommodity flow of value q in (D, c, π) , or
- (2) a partition (U, \bar{U}) of $V(D)$ of weighted ratio cost $\mathcal{O}(q \log p)$, where $p := |\{v \in V(D) : \pi(v) > 0\}|$.

We remark that the formulation of [14, Theorem 17] is non-algorithmic, but the proof is effective and can be turned into a randomized polynomial-time algorithm; see the remarks after [14, Theorem 7].

Similarly as in [11], for our purposes we need to translate the duality provided by **Theorem 6** to the setting of vertex separators, vertex-capacitated flows, and (vertex-)weighted undirected graphs. The translation uses a standard, simple reduction scheme: splitting every vertex u into two. However, we need to first introduce the relevant terminology.

The role of multicommodity flows will be played by *concurrent flows*. For a weighted graph G , a *concurrent flow* in G is a function $\lambda : \mathcal{P}(G) \rightarrow \mathbb{R}_{\geq 0}$ that assigns each path $P \in \mathcal{P}(G)$ the amount of flow passing through P , so that for every pair of vertices $(u, v) \in V(G) \times V(G)$ with $u \neq v$, we have

$$\sum_{P \in \mathcal{P}_{u,v}(G)} \lambda(P) = \mathbf{w}(u)\mathbf{w}(v).$$

By $\gamma_{G,\lambda}(v)$ we denote the total amount of flow passing through v , that is,

$$\gamma_{G,\lambda}(v) := \sum_{P \in \mathcal{P}(G), v \in P} \lambda(P) \tag{1}$$

The *congestion* γ of the concurrent flow λ is $\max_{v \in V(G)} \gamma_{G,\lambda}(v)$.

On the other hand, the role of balanced cuts will be played by balanced separations. Recall that a *separation* in a (weighted) graph G is a pair of vertex subsets (A, B) such that $A \cup B = V(G)$ and there is no edge with one endpoint in $A \setminus B$ and the other in $B \setminus A$. Then $A \cap B$ is the *separator* of the separation (A, B) . Assuming G is weighted, we define the *sparsity* of (A, B) as

$$\alpha_G(A, B) := \frac{|A \cap B|}{\mathbf{w}(A) \cdot \mathbf{w}(B)}.$$

With all the definitions in place, we can prove the desired variant of the result of Leighton and Rao.

Lemma 3 (undirected vertex-weighted variant of [14, Theorem 17]). *There is a randomized polynomial-time algorithm that, given a weighted n -vertex graph G and a real $\gamma > 0$, outputs either:*

- (1) a concurrent flow in G with congestion at most γ , or
- (2) a separation in G of sparsity $\mathcal{O}\left(\frac{\log n}{\gamma}\right)$.

241 *Proof.* We perform the standard trick of splitting every vertex into two. Let D be the directed graph
 242 with edge capacities $c : E(D) \rightarrow \mathbb{R}_{\geq 0}$ obtained from G as follows:

- 243 • for every $u \in V(G)$, add to D two vertices u_{in} and u_{out} and connect them by an edge $(u_{\text{in}}, u_{\text{out}})$ of
 244 capacity 1 and an edge $(u_{\text{out}}, u_{\text{in}})$ of capacity ∞ ; and
- 245 • for every edge $uv \in E(G)$, add to D edges $(u_{\text{out}}, v_{\text{in}})$ and $(v_{\text{out}}, u_{\text{in}})$ of capacity ∞ .

246 Edges of capacity ∞ can be avoided by using a large constant, e.g. $|V(G)|^2 \cdot \mathbf{w}(G) + 1$, instead. Finally,
 247 for every $u \in V(G)$, we set

$$248 \quad \pi(u_{\text{in}}) := \mathbf{w}(u) \quad \text{and} \quad \pi(u_{\text{out}}) := 0.$$

249 Thus, (D, c, π) is an instance of the Directed Product Multicommodity Flow Problem. Let us verify that
 250 multicommodity flows in this instance are in correspondence to concurrent flows in G .

251 **Claim 1.** *For any real $q > 0$, (D, c, π) has a multicommodity flow of value q if and only if G has
 252 a concurrent flow of congestion $1/q$.*

253 *Proof.* First, suppose there is a multicommodity flow f in (D, c, π) that, for every pair of distinct vertices
 254 $u, v \in V(G)$, sends $q \cdot \pi(u_{\text{in}})\pi(v_{\text{in}})$ units of flow from u_{in} to v_{in} . Note that there is a natural map from
 255 $\mathcal{P}_{u_{\text{in}}, v_{\text{in}}}(D)$ to $\mathcal{P}_{u, v}(G)$ defined by contracting every vertex pair $w_{\text{in}}, w_{\text{out}}$ back to w . Hence, we may map
 256 f to a flow λ_0 in G that sends $q \cdot \mathbf{w}(u)\mathbf{w}(v)$ units of flow from u to v , for every pair of distinct vertices
 257 u, v , and has congestion 1 (this corresponds to the edges of the form $(w_{\text{in}}, w_{\text{out}})$ having capacity 1). By
 258 rescaling λ_0 by a multiplicative factor of $1/q$, we obtain a concurrent flow λ in G with congestion $1/q$.

259 Conversely, suppose G has a concurrent flow λ of congestion $1/q$. Similarly as above, we may map
 260 λ to a multicommodity flow f_0 in (D, c, π) that sends $\mathbf{w}(u)\mathbf{w}(v) = \pi(u_{\text{in}})\pi(v_{\text{in}})$ units of flow from u_{in}
 261 to v_{in} , for every pair of distinct vertices u, v , and uses capacity $1/q$ at every edge of the form $(w_{\text{in}}, w_{\text{out}})$.
 262 By rescaling f_0 by a multiplicative factor of q , we obtain a multicommodity flow f of value q that
 263 respects all the edge capacities. \square

264 We apply the algorithm of [Theorem 6](#) to (D, c, π) with $q = 1/\gamma$. If this application yields a multi-
 265 commodity flow of value $1/\gamma$, then by [Claim 1](#) we may obtain a concurrent flow in G with congestion γ ,
 266 which is outcome [\(1\)](#). So we may assume that we have computed a partition (U, \bar{U}) of $V(D)$ with
 267 weighted ratio cost bounded by $\mathcal{O}(\log n/\gamma)$.

268 Observe that every edge in $E_D(U, \bar{U})$ must be of the form $(u_{\text{in}}, u_{\text{out}})$ for some $u \in V(G)$, because
 269 the other edges have infinite capacity, which would render the weighted ratio cost of (U, \bar{U}) infinite (or
 270 too large). Let $S \subseteq V(G)$ be the set comprised of all vertices u such that $(u_{\text{in}}, u_{\text{out}}) \in E_D(U, \bar{U})$. Then,

$$271 \quad \sum_{e \in E_D(U, \bar{U})} c(e) = |E_D(U, \bar{U})| = |S|.$$

272 Let $A \subseteq V(G)$ be the set of vertices u such that $u_{\text{in}} \in U$, and $B \subseteq V(G)$ be the set of vertices u such
 273 that $u_{\text{out}} \in \bar{U}$; thus $S = A \cap B$. Note that due to the existence of an edge $(u_{\text{out}}, u_{\text{in}})$ of capacity ∞ , for
 274 each vertex $u \in V(G)$ we must have $u \in A$ or $u \in B$. Moreover, since every edge uv of G gives rise
 275 to edges $(u_{\text{out}}, v_{\text{in}})$ and $(v_{\text{out}}, u_{\text{in}})$ of capacity ∞ , it cannot happen that $u \in A \setminus B$ (implying $u_{\text{out}} \in U$)
 276 and $v \in B \setminus A$ (implying $v_{\text{in}} \in \bar{U}$) simultaneously. We conclude that (A, B) is a separation of G , with
 277 separator $S = A \cap B$.

278 Finally, observe that the sparsity of (A, B) is

$$279 \quad \alpha_G(A, B) = \frac{|S|}{\mathbf{w}(A) \cdot \mathbf{w}(B)} = \frac{\sum_{e \in E_D(U, \bar{U})} c(e)}{\mathbf{w}(A) \cdot \mathbf{w}(B)}.$$

Note that $\pi(U) = \mathbf{w}(A)$ and $\pi(\bar{U}) = \mathbf{w}(B \setminus A) \leq \mathbf{w}(B)$. Thus we have

$$\alpha_G(A, B) \leq \frac{\sum_{e \in E_D(U, \bar{U})} c(e)}{\pi(U) \cdot \pi(\bar{U})},$$

which in turn is bounded by $\mathcal{O}(\log n / \gamma)$. So the separation (A, B) constitutes a valid outcome (2). \square

We now transform, by repeated applications, the second outcome of Lemma 3 into a balanced separator of the graph. This is analogous to [11, Lemma 4.1], but we repeat the argument here for the convenience of the reader, and also because the presentation of [11] claims a slightly stronger result assuming a faulty lemma from [13].

Lemma 4. *There is a randomized polynomial-time algorithm that, given a weighted n -vertex graph G with $W := \mathbf{w}(G)$ and a real $\gamma > 0$, outputs either:*

- (1) *an induced subgraph G' of G such that $\mathbf{w}(G') \geq W/2$, a concurrent flow in G' with congestion at most γ , and a subset $S \subseteq V(G)$ of size $\mathcal{O}\left(\frac{W^2 \log n}{\gamma}\right)$ such that $S \cap V(G') = \emptyset$ and every connected component of $G - S$ is either disjoint from G' or contained in G' ; or*
- (2) *a balanced separator of G of size $\mathcal{O}\left(\frac{W^2 \log n}{\gamma}\right)$.*

Proof. We build a sequence G_1, G_2, G_3, \dots of induced subgraphs of G as follows. Note here that each G_i , as an induced subgraph of G , inherits the weight function $\mathbf{w} := \mathbf{w}_G$ from G .

Set $G_1 := G$ and $i \leftarrow 1$. Perform the following loop: While $\mathbf{w}(G_i) \geq W/2$, apply Lemma 3 to G_i with congestion parameter γ . If this application yields outcome (1), then we break out of the while loop. Otherwise, we obtain a separation (A_i, B_i) in G_i with sparsity at most $\frac{c \log n}{\gamma}$, for some universal constant c . By swapping A_i and B_i if necessary, we assume that $\mathbf{w}(A_i) \geq \mathbf{w}(B_i)$. We set $G_{i+1} := G_i[A_i \setminus B_i]$ and $i \leftarrow i + 1$, and proceed to the next iteration. This ends the body of the loop.

When we exit the while loop, we have built separations $(A_1, B_1), \dots, (A_\ell, B_\ell)$ of induced subgraphs $G_1 \supseteq \dots \supseteq G_\ell$, respectively. For $i \in [\ell]$, denote $S_i := A_i \cap B_i$. We claim that if we have not broken out of the while loop (upon finding a concurrent flow), then $S := \bigcup_{i \in [\ell]} S_i$ is a balanced separator of G . Note that by construction, $V(G) \setminus S$ is partitioned into $B_1 \setminus A_1, B_2 \setminus A_2, \dots, B_\ell \setminus A_\ell$, and $A_\ell \setminus B_\ell$, with no edge between two distinct parts. For each $i \in [\ell]$, we have $\mathbf{w}(B_i \setminus A_i) \leq W/2$, because $\mathbf{w}(B_i) \leq \mathbf{w}(A_i)$. By the condition of exiting the loop, we have $\mathbf{w}(A_\ell \setminus B_\ell) \leq W/2$. So S is indeed a balanced separator of G .

We are left with arguing that $|S| \leq \mathcal{O}\left(\frac{W^2 \log n}{\gamma}\right)$. For every $i \in [\ell]$ we have

$$\alpha_{G_i}(A_i, B_i) = \frac{|S_i|}{\mathbf{w}(A_i) \cdot \mathbf{w}(B_i)} \leq \frac{c \log n}{\gamma}, \text{ thus}$$

$$|S| = \sum_{i \in [\ell]} |S_i| \leq \frac{c \log n}{\gamma} \cdot \sum_{i \in [\ell]} \mathbf{w}(A_i) \cdot \mathbf{w}(B_i).$$

Note that the sets B_1, \dots, B_ℓ are pairwise disjoint, and we can upper bound every term $\mathbf{w}(A_i)$ by the total weight W . Therefore

$$|S| \leq \frac{cW \log n}{\gamma} \cdot \sum_{i \in [\ell]} \mathbf{w}(B_i) \leq \frac{cW^2 \log n}{\gamma},$$

and S constitutes a valid outcome ((2)).

Now, suppose we have broken out of the while loop because a concurrent flow λ of congestion at most γ was found in $G' := G[A_\ell \setminus B_\ell]$. By the invariant of the loop, G' is an induced subgraph of G such that $\mathbf{w}(G') \geq W/2$. And by construction, S is disjoint with $V(G')$ and every connected component of $G - S$ is either contained in G' or disjoint with G' . Also, by the same calculation as above, we have $|S| \leq \mathcal{O}\left(\frac{W^2 \log n}{\gamma}\right)$. We conclude that G', λ , and S satisfy the desired properties of outcome ((1)). \square

3 Proof of **Theorem 2**

In this section we prove our main result, **Theorem 2**. By the following result, the statement can be reduced to the $d = 3$ case. Recall here that G^d is the d th power of G — the graph obtained from G by making two vertices of $V(G)$ adjacent if and only if they are at distance at most d in G .

Theorem 7 ([7, Theorem 3]). *If G excludes H as a d -fat minor, then G^d excludes H as a 3-fat minor.*

Indeed, assuming $d > 3$, we apply the 3-fat minor case of **Theorem 2** to G^d . This yields either a 3-fat model of H in G^d , which can be turned into a d -fat model of H in G (the proof of **Theorem 7** yields a polynomial-time algorithm), or a set $S \subseteq V(G^d) = V(G)$ that is a balanced separator of G^d and is $(\mathcal{O}(\|H\|^2 \cdot n^{1/2+\varepsilon}), \mathcal{O}(1/\varepsilon))$ -coverable in G^d . Clearly, as G is a subgraph of G^d , S is also a balanced separator of G . Moreover, as the distances in G are at most d times larger than in G^d , it follows that S is $(\mathcal{O}(\|H\|^2 \cdot n^{1/2+\varepsilon}), \mathcal{O}(d/\varepsilon))$ -coverable in G .

Therefore, we are left with proving the following statement.

Theorem 8. *For every graph H and real $\varepsilon > 0$, every n -vertex weighted graph that excludes H as a 3-fat minor has an $(\mathcal{O}(\|H\|^2 \cdot n^{1/2+\varepsilon}), \mathcal{O}(1/\varepsilon))$ -coverable balanced separator.*

Furthermore, there is a randomized polynomial-time algorithm that, given a weighted graph G , either returns such a balanced separator in G or finds a 3-fat model of H in G .

Proof. We focus on showing the algorithmic claim, as it implies the first part of the theorem statement. By adding some edges to H if necessary, we may assume that H has no isolated vertices; this increases $\|H\|$ only by a constant multiplicative factor, and if a 3-fat minor model of this supergraph is found, one can effectively restrict the model to one for the original H . Denote $W := \mathbf{w}_G(G)$, for brevity.

By applying the algorithm of **Lemma 1**, we may compute, in randomized polynomial time, a connected partition \mathcal{X} of G of strong diameter at most $32/\varepsilon$ such that for every vertex v of G , the ball $B_G(v, 2)$ intersects at most $\mathcal{O}(n^\varepsilon)$ clusters of \mathcal{X} . We observe that from this it follows that there are relatively few clusters of \mathcal{X} that are close to each other. Formally, we consider the set

$$\Xi := \{ (X, X') \in \mathcal{X} \times \mathcal{X} \mid \text{dist}_G(X, X') \leq 2 \}.$$

Claim 2. $|\Xi| \leq \mathcal{O}(n^{1+\varepsilon})$.

Proof. Since every vertex of G is at distance at most 2 from at most $\mathcal{O}(n^\varepsilon)$ clusters of \mathcal{X} , each cluster $X \in \mathcal{X}$ is at distance at most 2 from at most $|X| \cdot \mathcal{O}(n^\varepsilon)$ other clusters of \mathcal{X} . Thus, we have

$$|\Xi| \leq \sum_{X \in \mathcal{X}} (|X| \cdot \mathcal{O}(n^\varepsilon)) = \left(\sum_{X \in \mathcal{X}} |X| \right) \cdot \mathcal{O}(n^\varepsilon) = n \cdot \mathcal{O}(n^\varepsilon) = \mathcal{O}(n^{1+\varepsilon}). \quad \square$$

Let c be the constant hidden in the $\mathcal{O}(\cdot)$ notation of **Claim 2**; that is, we have $|\Xi| \leq cn^{1+\varepsilon}$. We let $\widehat{G} := G/\mathcal{X}$ be the quotient graph. Recall that \widehat{G} is weighted with the weight function defined as $\mathbf{w}_{\widehat{G}}(X) := \mathbf{w}_G(X)$ for all $X \in \mathcal{X}$, so in particular $\mathbf{w}_{\widehat{G}}(\widehat{G}) = \mathbf{w}_G(G) = W$.

We apply the algorithm of **Lemma 4** to the (weighted) graph \widehat{G} with

$$\gamma := \frac{W^2}{32h\sqrt{c} \cdot n^{\frac{1+\varepsilon}{2}}}$$

where $h := \|\widehat{H}\| = |V(\widehat{H})| + |E(\widehat{H})|$ is the total number of vertices and edges of \widehat{H} ; note that $h \leq 5\|H\|$. If this application yields outcome (2), we have obtained a balanced separator $\mathcal{S} \subseteq \mathcal{X}$ of \widehat{G} of size $\mathcal{O}(W^2 \log n / \gamma) \leq \mathcal{O}(\|H\| \cdot n^{1/2+\varepsilon})$. Hence, $\bigcup \mathcal{S}$ is a balanced separator of G that can be covered by $|\mathcal{S}| = \mathcal{O}(\|H\| \cdot n^{1/2+\varepsilon})$ balls of radius $32/\varepsilon$ in G .

We thus assume that the application of **Lemma 4** yields outcome (1). That is, we have obtained an induced subgraph \widehat{G}' of \widehat{G} with $\mathbf{w}_{\widehat{G}'}(\widehat{G}') \geq \mathbf{w}_{\widehat{G}}(\widehat{G})/2 = W/2$, a concurrent flow λ on \widehat{G}' with congestion

at most γ , and a subset $\widehat{S} \subseteq V(\widehat{G}) = \mathcal{X}$ of size at most $\mathcal{O}(W^2 \log n / \gamma) = \mathcal{O}(\|H\| \cdot n^{1/2+\varepsilon})$ such that every connected component of $\widehat{G} - \widehat{S}$ is disjoint from \widehat{G}' or contained in \widehat{G}' . Note that every component of $\widehat{G} - \widehat{S}$ disjoint from \widehat{G}' must have weight at most $W/2$, for \widehat{G}' has weight at least $W/2$.

We now partition the vertex set of \widehat{G}' into two sets \mathcal{Z} and \mathcal{U} as follows:

- \mathcal{Z} comprises all $X \in V(\widehat{G}')$ such that $\mathbf{w}_G(X) \geq \frac{W}{4h^2}$; and
- \mathcal{U} comprises all the other vertices of \widehat{G}' , i.e., those $X \in V(\widehat{G}')$ that satisfy $\mathbf{w}_G(X) < \frac{W}{4h^2}$.

Note that since $\mathbf{w}_G(\widehat{G}') \leq W$, we have $|\mathcal{Z}| \leq 4h^2 \leq \mathcal{O}(\|H\|^2)$.

We first consider the case when $\mathbf{w}_{\widehat{G}}(\mathcal{Z}) \geq \mathbf{w}_{\widehat{G}}(\mathcal{U})$. This means that \mathcal{Z} contains at least half of the weight of \widehat{G}' , and hence $\widehat{S} \cup \mathcal{Z}$ is a balanced separator of \widehat{G} . Therefore, $S := \bigcup(\widehat{S} \cup \mathcal{Z})$ is a balanced separator of G , and it can be covered by $|\widehat{S}| + |\mathcal{Z}| \leq \mathcal{O}(\|H\|^2 \cdot n^{1/2+\varepsilon})$ balls of radius $32/\varepsilon$ in G .

Hence, from now on we assume that $\mathbf{w}_{\widehat{G}}(\mathcal{U}) > \mathbf{w}_{\widehat{G}}(\mathcal{Z})$. In particular, this implies that

$$M := \mathbf{w}_{\widehat{G}}(\mathcal{U}) \geq \mathbf{w}_{\widehat{G}}(\widehat{G}')/2 \geq W/4.$$

We now construct a crude 3-fat model $(\phi : V(\dot{H}) \rightarrow V(\widehat{G}), \pi : E(\dot{H}) \rightarrow \mathcal{P}(\widehat{G}))$ of \dot{H} in \widehat{G} as follows. Let \mathcal{D} be the probability distribution over \mathcal{U} that assigns probability $\mathbf{w}_G(X)/M$ to every $X \in \mathcal{U}$. For every pair (X, Y) of distinct elements of \mathcal{U} , let $\mathcal{D}_{X,Y}$ be the probability distribution over $\mathcal{P}_{X,Y}(\widehat{G}')$ that assigns probability

$$\frac{\lambda(P)}{\mathbf{w}_G(X)\mathbf{w}_G(Y)} \text{ to every } P \in \mathcal{P}_{X,Y}(\widehat{G}').$$

For every $u \in V(\dot{H})$, we sample $\phi(u)$ independently from \mathcal{D} . Then, for every $uv \in E(\dot{H})$, if $\phi(u) \neq \phi(v)$, then we sample $\pi(uv)$ independently from $\mathcal{D}_{\phi(u),\phi(v)}$. If $\phi(u) = \phi(v)$, then $\pi(uv)$ is instead set to the trivial $\phi(u)$ - $\phi(v)$ path.

We consider the following two events:

I : The mapping ϕ is injective.

F : For every pair of edges $uv, u'v' \in E(\dot{H})$ with $|\{u, v, u', v'\}| = 4$, we have $\text{dist}_{\widehat{G}}(\pi(uv), \pi(u'v')) \geq 3$.

As usual, by \bar{I} and \bar{F} we denote the complements of those events, respectively.

We first bound the probability of \bar{I} :

$$\begin{aligned} \mathbb{P}[\bar{I}] &\leq \sum_{u,v \in V(\dot{H}), u \neq v} \mathbb{P}[\phi(u) = \phi(v)] \\ &= \sum_{u,v \in V(\dot{H}), u \neq v} \sum_{X \in \mathcal{U}} \mathbb{P}[\phi(u) = X] \cdot \mathbb{P}[\phi(v) = X] \\ &= \sum_{u,v \in V(\dot{H}), u \neq v} \sum_{X \in \mathcal{U}} \frac{\mathbf{w}_G(X)^2}{M^2} \\ &\leq \binom{h}{2} \cdot \frac{W}{4h^2} \cdot \frac{1}{M^2} \cdot \sum_{X \in \mathcal{U}} \mathbf{w}_G(X) && \text{by } \mathbf{w}_G(X) \leq \frac{W}{4h^2} \text{ for each } X \in \mathcal{U}, \\ &\leq \binom{h}{2} \cdot \frac{W}{4Mh^2} \leq \frac{1}{2} && \text{by } \sum_{X \in \mathcal{U}} \mathbf{w}_G(X) = M \geq W/4. \end{aligned} \quad (2)$$

We proceed to bounding the probability of \bar{F} . For a fixed $X \in V(\widehat{G}')$ and $uv \in E(\dot{H})$ consider the event $A_{X,uv}$ defined as follows: $\phi(u) \neq \phi(v)$ and $X \in \pi(uv)$. We first bound the probability of $A_{X,uv}$:

$$\mathbb{P}[A_{X,uv}] = \sum_{(Y,Z) \in \mathcal{U} \times \mathcal{U}, Y \neq Z} \frac{\mathbf{w}_G(Y)\mathbf{w}_G(Z)}{M^2} \sum_{P \in \mathcal{P}_{Y,Z}(\widehat{G}'), X \in P} \frac{\lambda(P)}{\mathbf{w}_G(Y)\mathbf{w}_G(Z)}$$

$$\begin{aligned}
&= \frac{1}{M^2} \sum_{(Y,Z) \in \mathcal{U} \times \mathcal{U}, Y \neq Z} \frac{\mathbf{w}_G(Y)\mathbf{w}_G(Z)}{\mathbf{w}_G(Y)\mathbf{w}_G(Z)} \sum_{P \in \mathcal{P}_{Y,Z}(\widehat{G}'), X \in P} \lambda(P) \\
&= \frac{1}{M^2} \sum_{P \in \mathcal{P}(\widehat{G}'), X \in P} \lambda(P) = \frac{\gamma_{\widehat{G}', \lambda}(X)}{M^2} \leq \frac{\gamma}{M^2}.
\end{aligned}$$

Let us define the set of relevant pairs of edges of \ddot{H} :

$$\Pi := \{(uv, u'v') \in E(\ddot{H}) \times E(\ddot{H}) \mid |\{u, v, u', v'\}| = 4\}.$$

Note that for any $(uv, u'v') \in \Pi$, $(\phi(u), \phi(v), \pi(uv))$ and $(\phi(u'), \phi(v'), \pi(u'v'))$ are independent joint variables, and hence $A_{X,uv}$ and $A_{X',u'v'}$ are also independent events. Therefore, for all $X, X' \in V(\widehat{G}')$,

$$\mathbb{P}[A_{X,uv} \cap A_{X',u'v'}] = \mathbb{P}[A_{X,uv}] \cdot \mathbb{P}[A_{X',u'v'}] \leq \frac{\gamma^2}{M^4} \leq 256 \cdot \frac{\gamma^2}{W^4} = \frac{1}{4h^2c \cdot n^{1+\varepsilon}}. \quad (3)$$

Observe that assuming ϕ is injective (that is, event I), that (ϕ, π) is a crude 3-fat model of \ddot{H} is equivalent to the following assertion: for all $(uv, u'v') \in \Pi$ and $(X, X') \in \Xi$, we do not have $X \in \pi(uv)$ and $X' \in \pi(u'v')$ simultaneously. Therefore, we have the following inclusion of events:

$$\overline{F} \cap I \subseteq \bigcup_{(uv, u'v') \in \Pi} \bigcup_{(X, X') \in \Xi} A_{X,uv} \cap A_{X',u'v'}.$$

Hence, we get

$$\begin{aligned}
\mathbb{P}[\overline{F} \cap I] &\leq |\Pi| \cdot |\Xi| \cdot \frac{1}{4h^2c \cdot n^{1+\varepsilon}} && \text{by union bound and (3),} \\
&\leq \frac{1}{4} && \text{by Claim 2 and } |\Pi| \leq h^2. \quad (4)
\end{aligned}$$

By combining (2) and (4) we conclude that $\mathbb{P}[F] \geq \mathbb{P}[F \cap I] \geq \frac{1}{4}$; or in other words, (ϕ, π) is a crude 3-fat model of \ddot{H} in \widehat{G} with probability at least $\frac{1}{4}$. By Lemma 2, this model can be turned into a 3-fat model $(\{\widehat{B}_u\}_{u \in V(H)}, \{\widehat{B}_e\}_{e \in E(H)})$ of H . Finally, since the distances between the elements of \mathcal{X} in \widehat{G} are not larger than in G , we conclude that taking

$$B_f := G \left[\bigcup_{X \in V(\widehat{B}_f)} X \right] \quad \text{for all } f \in V(H) \cup E(H)$$

yields a 3-fat model of H in G .

Taking into account the probability that any of the (randomized) algorithms of Lemmas 1 and 3 returns an incorrect output, which can be bounded by $\frac{1}{8}$, we conclude that the presented algorithm finds either an $(\mathcal{O}(\|H\|^2 \cdot n^{\frac{1}{2}+\varepsilon}), \mathcal{O}(1/\varepsilon))$ -coverable balanced separator of G or a 3-fat model of H in G with probability at least $\frac{1}{8}$. As usual, the probability of success can be boosted to $1 - \frac{1}{2^{n^{\mathcal{O}(1)}}$ by making $n^{\mathcal{O}(1)}$ independent repetitions of the algorithm. \square

4 Proofs of Theorems 3 and 4

In this section we prove our results for graphs excluding an induced minor: Theorems 3 and 4. We need some terminology. A *star partition* of a graph G is a connected partition \mathcal{X} of G such that for every $X \in \mathcal{X}$, there is a vertex $v_X \in X$ satisfying $X \subseteq N[v_X]$. In other words, it is a connected partition of strong radius at most 1.

Lemma 5. *Every n -vertex graph G admits a star partition \mathcal{X} such that G/\mathcal{X} has $\mathcal{O}(n^{4/3} \log^{2/3} n)$ edges.*

426 *Proof.* We set $k := \lceil n^{1/3} \log^{2/3} n \rceil$, and assume that $k \geq 2$ without loss of generality. Let v_1, \dots, v_h
427 be a maximal sequence of distinct vertices such that for every $i \in [h]$, v_i has at most k neighbors in
428 $G - \{v_1, \dots, v_{i-1}\}$. Let $M := \{v_1, \dots, v_h\}$, and let E_M be the set of edges of G with at least one endpoint
429 in M . By construction, $|E_M| \leq k|M| \leq kn$.

430 By the maximality of M , the graph $G' := G - M$ has minimum degree greater than k . Thus G'
431 admits a dominating set of size at most $\frac{2(n-h)\log k}{k}$ [5]. In particular, G' admits a star partition \mathcal{X}' with
432 at most $\frac{2(n-h)\log k}{k}$ parts. We extend it to a star partition \mathcal{X} of G by adding every singleton $\{v_1\}, \dots, \{v_h\}$.

433 The number of edges of G/\mathcal{X} is at most

$$434 |E_M| + \binom{\frac{2(n-h)\log k}{k}}{2} \leq kn + \frac{4n^2 \log^2 k}{k^2} = \mathcal{O}(n^{4/3} \log^{2/3} n). \quad \square$$

435 Then **Theorem 3** is an easy consequence of combining **Lemma 5** and **Theorem 5**.

436 *Proof of Theorem 3.* Apply **Lemma 5** to G , yielding a star partition \mathcal{X} of G such that the quotient graph
437 G/\mathcal{X} has $\mathcal{O}(n^{4/3} \log^{2/3} n)$ edges. Noting that G/\mathcal{X} also excludes H as an induced minor, we may
438 apply **Theorem 5** to G/\mathcal{X} , thus obtaining a balanced separator \mathcal{S} of G/\mathcal{X} of size $\mathcal{O}_H(\sqrt{m} \log m) =$
439 $\mathcal{O}_H(n^{2/3} \log^{4/3} n)$. Then $\bigcup \mathcal{S}$ is a balanced separator in G that is $(\mathcal{O}_H(n^{2/3} \log^{4/3} n), 1)$ -coverable. \square

440 For **Theorem 4**, we use the following result of Chudnovsky et al. [6] together with a classic conse-
441 quence of a result of Kühn and Osthus [12].

442 **Theorem 9** ([6, Theorem 9.1]). *Let G be a $K_{t,t}$ -induced-minor-free graph, for some $t \in \mathbb{N}$. Then G admits*
443 *a connected partition \mathcal{X} in which every part is $(1, 4)$ -coverable and such that the graph G/\mathcal{X} has chromatic*
444 *number at most $t \cdot 2^t$.*

445 **Theorem 10** (follows from [12, Theorem 1]). *For every $s \in \mathbb{N}$, every n -vertex graph G that excludes K_s*
446 *as a subgraph and $K_{s,s}$ as an induced minor has $\mathcal{O}_s(n)$ edges.*

447 As observed in [6, Proposition 10.3], the combination of these two results immediately yields the
448 following. We note that in [6], they use more recent, stronger results instead of the standard theorem
449 of Kühn and Osthus [12], but the latter suffices for our purpose.

450 **Lemma 6.** *For every $t \in \mathbb{N}$, every n -vertex $K_{t,t}$ -induced-minor-free graph G has a connected partition \mathcal{X}*
451 *such that every part of \mathcal{X} is $(1, 4)$ -coverable and the quotient graph G/\mathcal{X} has $\mathcal{O}_t(n)$ edges.*

452 Now **Theorem 4** follows from combining **Lemma 6** and **Theorem 5** in exactly the same manner as in
453 the proof of **Theorem 3**.

454 5 Conjectures

455 Our work raises multiple questions about the existence of coverable balanced separators in fat-minor-
456 free and induced-minor-free graphs. The most prominent one is whether **Conjecture 1** holds, or at least
457 improving the additional $\mathcal{O}(n^\epsilon)$ factor to a polylogarithmic function of n . The only source of the $\mathcal{O}(n^\epsilon)$
458 factor is the clustering scheme of **Lemma 1**, hence we propose the following.

459 **Conjecture 2.** *For every graph H there are constants $p, d \in \mathbb{N}$ such that every graph G that excludes H*
460 *as a 3-fat minor admits a connected partition \mathcal{X} of weak diameter at most d such that for every $u \in V(G)$,*
461 *$B_G(u, 2)$ intersects at most p clusters of \mathcal{X} .*

462 We note that resolving **Conjecture 2** would not automatically resolve **Conjecture 1** because of
463 the additional $\log n$ factor stemming from the gap between multicommodity flows and balanced cuts
464 (**Theorem 6**). However, proving a weaker form of **Conjecture 2** with p replaced by $\mathcal{O}(\log^\alpha n)$, for some
465 $\alpha \in \mathbb{R}_{\geq 0}$, would directly translate to the weaker form of **Conjecture 1** with c replaced by $\mathcal{O}(\log^{\alpha+1} n)$.

We note that in [Conjecture 2](#), it is not at all clear that the condition “every ball of radius 2 intersects few parts of \mathcal{X} ” is the only correct one; other forms of asserting the sparsity of \mathcal{X} may also be useful.

Noting that containing H as a 2-fat minor entails containing H as an induced minor, [Theorem 2](#) implies that any n -vertex H -induced-minor-free graph has an $(\mathcal{O}_H(n^{1/2+\varepsilon}), \mathcal{O}(1/\varepsilon))$ -coverable balanced separator, for any $\varepsilon > 0$. As witnessed by [Theorems 3](#) and [4](#), the induced-minor-free case may permit a broader range of tools, hence we state the following weaker form of [Conjecture 1](#).

Conjecture 3. *For every graph H there exist constants $c, r \in \mathbb{N}$ such that every n -vertex weighted graph G that excludes H as an induced minor has a balanced separator that is $(c\sqrt{n}, r)$ -coverable.*

It would be particularly elegant if $r = 1$ sufficed in [Conjecture 3](#), and [Theorem 3](#) gives an indication that this could indeed be the case. Towards this goal, we propose the following strengthening of [Lemma 5](#) as a first step.

Conjecture 4. *For every graph H there exist constants $c, q \in \mathbb{N}$ such that every n -vertex graph G that excludes H as an induced minor admits a connected partition \mathcal{X} such that every cluster of \mathcal{X} is $(q, 1)$ -coverable and the quotient graph G/\mathcal{X} has at most $c \cdot n$ edges.*

Again, a weaker form of this statement, with c or q (or both) replaced by a subpolynomial function of n , would also be interesting.

Acknowledgements. This work was initiated at Dagstuhl Seminar 25212: Metric Sketching and Dynamic Algorithms for Geometric and Topological Graphs. We thank the organizers and the other participants for creating a wonderfully friendly and productive atmosphere.

References

- [1] Tara Abrishami, Jadwiga Czyżewska, Kacper Kluk, Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzażewski. On coarse tree decompositions and coarse balanced separators. *ArXiv preprint*, abs/2502.20182, 2025.
- [2] Sandra Albrechtsen and James Davies. Counterexample to the conjectured coarse grid theorem. *ArXiv preprint*, abs/2508.15342, 2025.
- [3] Sandra Albrechtsen, Marc Distel, and Agelos Georgakopoulos. Small counterexamples to the fat minor conjecture. *ArXiv preprint*, abs/2601.05761, 2026.
- [4] Noga Alon, Paul Seymour, and Robin Thomas. A separator theorem for nonplanar graphs. *Journal of the American Mathematical Society*, 3(4):801–808, 1990.
- [5] Vladimir I. Arnautov. Estimation of the exterior stability number of a graph by means of the minimal degree of the vertices. *Prikladnaya Matematika i Programirovanie*, 11(3-8):126, 1974.
- [6] Maria Chudnovsky, Julien Codsı, Ajaykrishnan E. S., and Daniel Lokshantov. Induced minors and coarse tree decompositions. *ArXiv preprint*, abs/2603.11379, 2026.
- [7] James Davies, Robert Hickingbotham, Freddie Illingworth, and Rose McCarty. Fat minors cannot be thinned (by quasi-isometries). *ArXiv preprint*, abs/2405.09383, 2024.
- [8] Zdenek Dvořák and Sergey Norin. Treewidth of graphs with balanced separations. *Journal of Combinatorial Theory B*, 137:137–144, 2019.
- [9] Arnold Filtser. Scattering and sparse partitions, and their applications. *ACM Transactions on Algorithms*, 20(4):30:1–30:42, 2024.

- 505 [10] Agelos Georgakopoulos and Panos Papasoglu. Graph minors and metric spaces. *Combinatorica*,
506 45(3):33, 2025.
- 507 [11] Tuukka Korhonen and Daniel Lokshtanov. Induced-minor-free graphs: Separator theorem, subex-
508 ponential algorithms, and improved hardness of recognition. In *2024 Annual ACM-SIAM Symposium*
509 *on Discrete Algorithms, SODA 2024*, pages 5249–5275, 2024.
- 510 [12] Daniela Kühn and Deryk Osthus. Induced subdivisions in $K_{s,s}$ -free graphs of large average degree.
511 *Combinatorica*, 24(2):287–304, 2004.
- 512 [13] James R. Lee. Separators in region intersection graphs. In *8th Innovations in Theoretical Computer*
513 *Science Conference, ITCS 2017*, LIPIcs, pages 1:1–1:8. Schloss Dagstuhl – Leibniz-Zentrum für
514 Informatik, 2017.
- 515 [14] Frank Thomson Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their
516 use in designing approximation algorithms. *Journal of the ACM*, 46(6):787–832, 1999.
- 517 [15] Richard J. Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. *SIAM Journal*
518 *on Applied Mathematics*, 36(2):177–189, 1979.