Excluding $K_{2,t}$ as a fat minor
Working Draft

Sandra Albrechtsen[1], Marc Distel*[2], and Agelos Georgakopoulos†[3]

[1]Department of Mathematics, University of Hamburg, Bundesstrasse 55, 20146 Hamburg, Germany
[2]School of Mathematics, Monash University, Melbourne, Australia
[3]Mathematics Institute, University of Warwick, CV4 7AL, UK

September 5, 2025

Abstract

We prove that for every $t \in \mathbb{N}$, the graph $K_{2,t}$ satisfies the fat minor conjecture of Georgakopoulos and Papasoglu: for every $K \in \mathbb{N}$ there exist $M, A \in \mathbb{N}$ such that every graph with no K-fat $K_{2,t}$ minor is (M, A)-quasi-isometric to a graph with no $K_{2,t}$ minor. We use this to obtain an efficient algorithm for approximating the minimal multiplicative distortion of any embedding of a finite graph into a $K_{2,t}$-minor-free graph, answering a question of Chepoi, Dragan, Newman, Rabinovich, and Vaxès from 2012.

Keywords: coarse graph theory, quasi-isometry, asymptotic minor, approximation algorithm.

MSC 2020 Classification: 05C83, 05C10, 68R12, 05C63, 51F30.

1 Introduction

Coarse graph theory is a rapidly developing new area that studies graphs from a geometric perspective, and conversely, transfers graph-theoretic results to metric spaces. The focus is on large-scale properties of the graphs and spaces involved, in particular on properties that are stable under quasi-isometries (defined in Section 2.3). A central notion of this area is that of a *K-fat minor*, a geometric analogue of the classical notion of graph minor whereby branch sets are required to be at distance at least some distance K from each other, and the edges connecting them are replaced by long paths, also at distance K from each other, and from their non-incident branch sets; see Section 2.2 for details. We say that a graph J is an *asymptotic minor* of a graph G, if J is a K-fat minor of G for every $K \in \mathbb{N}$. For any fixed J, this property is easily seen to be invariant under quasi-isometry on G ([13, Observation 2.4]).

*Supported by Australian Government Research Training Program Scholarship.
†Supported by EPSRC grant EP/V009044/1.

Much of the impetus of coarse graph theory is due to the following conjecture of [13]:

Conjecture 1.1 ([13]). *For every finite graph J and every $K \in \mathbb{N}$ there exist $M, A \in \mathbb{N}$ such that every graph with no K-fat J minor is (M, A)-quasi-isometric to a graph with no J minor.*

In other words, the conjecture asks whether every graph (family) forbidding J as an asymptotic minor is (uniformly) quasi-isometric with a graph (family) forbidding J as a minor. This was a natural conjecture to make, as the converse is easily seen to be true. However, Conjecture 1.1 was disproven by Davies, Hickingbotham, Illingworth and McCarty [10]. In a companion paper [4] we will provide much smaller counterexamples; in particular, we will prove that it is false for $J = K_t, t \geq 6$, and for $K_{s,t}, s, t \geq 4$. Recently, Albrechtsen and Davies [2] also disproved a weaker version of Conjecture 1.1, stated in [10], postulating a quasi-isometry to a graph forbidding some possibly much larger graph J' as a minor.

This negative answer to Conjecture 1.1 fuels the interest in the broader quest, already initiated by Bonamy et al. [8], to understand graphs (or graph families) with a forbidden asymptotic minor. A substantial aspect of this quest, motivating the current paper, is to understand the limits of the validity of the conjecture. Several positive results have been obtained so far: Conjecture 1.1 is true e.g. for $J = K_3$ (more generally, for any cycle J) [13], for $J = K_{1,t}$ [13, 14], for $J = K_4^-$ [12, 6], $J = K_{2,3}$ [9, 12], and $J = K_4$ [6]. An important open question, due to its connection with induced minors, is whether Conjecture 1.1 is true for $K = 2$.

Given the above results, a central outstanding case towards understanding which graphs satisfy Conjecture 1.1 is the case $J = K_{2,t}, t \geq 4$. This question is implicit in earlier work of Chepoi, Dragan, Newman, Rabinovich and Vaxès [9], where a variant of the notion of fat minor is introduced. The aim of this paper is to settle this question in the affirmative; we prove

Theorem 1.2. *For every $t \in \mathbb{N}$ there exists a function $f : \mathbb{N} \to \mathbb{N}^2$ such that every graph with no K-fat $K_{2,t}$ minor is $f(K)$-quasi-isometric to a graph with no $K_{2,t}$ minor.*

We remark that this problem bears some similarity to the coarse Menger conjecture [5, 13], which has been disproven even in a much weaker form [15].

Our proof is constructive, and we obtain the bound $(9t^{12}K + 204t^9K, 1)$ on $f(K)$. In other words, the additive distortion we obtain is 1, and the multiplicative distortion $O(K)$. From this it is easy to obtain a map of additive distortion 0 (and still with multiplicative distortion $O(K)$ [13, Observation 2.2][1]).

Given a finite graph G, let $\alpha_t(G)$ denote the minimal multiplicative distortion of any embedding of G into a $K_{2,t}$-minor-free graph. Chepoi et al. [9] asked whether there is an efficient algorithm that approximates $\alpha_t(G)$ to a constant factor. Using the above remarks we answer this question in the affirmative:

[1] The additive error can always be hidden inside the multiplicative factor, unless more than one vertex of G is mapped to the same vertex of H. In this case, attach a star of size $|V(G)|$ to each vertex h of H (which does not create any $K_{2,t}$ minors), and for each vertex v of G previously mapped to h, map v to a distinct leaf of the star attached to h.

Corollary 1.3. *For every $t \in \mathbb{N}$, there is a polynomial-time algorithm that given a finite graph G, approximates $\alpha_t(G)$ up to a universal multiplicative constant.*

We prove this in Section 8, where we offer some related open problems.

1.1 Other problems

As mentioned above, Theorem 1.2 becomes false if we replace $K_{2,t}$ by $K_{4,t}, t \geq 4$ (even in a weak form as in Question 1.2 below), but we do not know if it is true for $K_{3,t}, t \geq 3$. The case $J = K_{3,3}$ is particularly important, as it is closely related to the 'coarse Kuratowski conjecture' of [13]:

Question 1.4. *Are there functions $f : \mathbb{N} \to \mathbb{N}^2$ and $s : \mathbb{N} \to \mathbb{N}$ such that every graph with no K-fat $K_{3,t}$ minor is $f(K)$-quasi-isometric to a graph with no $K_{3,s(t)}$ minor? Can we choose $s(t) = t$?*

Another question of [13] is for which J we can achieve $M = 1$ in Conjecture 1.1, and variants of this question are discussed by Nguyen, Scott and Seymour [14, 16]. Settling this for $J = K_{2,t}$ would be interesting, but our proof does not provide evidence.

1.2 Proof approach

Like many results in the area, our proof of Theorem 1.2 is achieved by decomposing the vertex set of the underlying graph G into 'bags', of bounded diameter, so that after collapsing each bag into a vertex, the resulting graph H is quasi-isometric to G. The standard technique is to achieve such a decomposition by first decomposing G into its distance layers from a fixed 'root' vertex, and place nearby vertices of a fixed layer, or a fixed number of consecutive layers, into a bag, see e.g. [13, Theorem 3.1]. Our decomposition is based on a rather intricate refinement of this technique, whereby the number of consecutive layers from which a bag is formed is not fixed but depends on the local structure. Once H is constructed, one then needs a way to turn any $K_{2,t}$ minor of H into a K-fat minor of G; this is not straightforward, one of the difficulties being that bags are not necessarily connected. Thus our proof requires new ideas involving a new way of forming branch sets in G out of bags in H by using vertices from bags of lower layers. To ensure that distinct branch sets are K-far apart, we use a new 'buffer zone' technique within each bag, i.e. a sequence of layers that can only be used to accommodate branch paths. A more detailed overview of our proof is given in Section 3.

2 Preliminaries

Graphs in this paper are allowed to be infinite, unless stated otherwise. We follow the basic graph-theoretic terminology of [11]; in particular, $\mathbb{N}$ includes 0, and we denote by $||G||$ the number of edges of a graph G. Note that if P is a path, then $||P||$ is its length. Moreover, a set U of vertices in a graph G is *connected*, if the subgraph $G[U]$ it induces is connected.

Given a graph G, we write $\mathcal{C}(G)$ for the set of components of G. Given a subgraph Y of G, the *boundary* $\partial_G Y$ of Y is the set of all vertices of Y that send an edge to $G - Y$. The *neighbourhood* $N_G(Y)$ of Y is the set of vertices of $G - Y$ sending an edge to Y (and therefore to $\partial_G Y$).

2.1 Distances

Let G be a graph. We write $d_G(v, u)$ for the distance between two vertices v and u in G. For two sets U and U' of vertices of G, we write $d_G(U, U')$ for the minimum distance of two elements of U and U', respectively. If one of U or U' is just a singleton, then we omit the braces, writing $d_G(v, U') := d_G(\{v\}, U')$ for $v \in V(G)$.

Given a set U of vertices of G, the *ball (in G) around U of radius $r \in \mathbb{N}$*, denoted by $B_G(U, r)$, is the set of all vertices in G of distance at most r from U in G. If $U = \{v\}$ for some $v \in V(G)$, then we again omit the braces, writing $B_G(v, r)$ instead of $B_G(\{v\}, r)$.

The *diameter* $\mathrm{diam}(G)$ of G is the smallest number $k \in \mathbb{N} \cup \{\infty\}$ such that $d_G(u, v) \leq k$ for every two $u, v \in V(G)$. If G is empty, then we define its diameter to be 0. We remark that if G is disconnected but not the empty graph, then its diameter is ∞. The *diameter of a set $U \subseteq V(G)$ in G*, denoted by $\mathrm{diam}_G(U)$, is the smallest number $k \in \mathbb{N}$ such that $d_G(u, v) \leq k$ for all $u, v \in U$ or ∞ if such a $k \in \mathbb{N}$ does not exist.

If Y is a subgraph of G, then we abbreviate $d_G(U, V(Y))$, $\mathrm{diam}_G(V(Y))$ and $B_G(V(Y), r)$ as $d_G(U, Y)$, $\mathrm{diam}_G(Y)$ and $B_G(Y, r)$, respectively.

Let G be a graph. We say that $U \subseteq V(G)$ is *K-near-connected* for $K \in \mathbb{N}$, if for every $x, y \in U$, there is a sequence $x = x_0, x_1, \ldots, x_k = y$ of vertices in U such that $d(x_i, x_{i+1}) \leq K$ for every $i < k$. Such a sequence $P = x_0, \ldots, x_k$ will be called an *K-near path* from x to y. A *K-near-component* of U is a maximal subset of U that is K-near-connected.

2.2 Fat minors

Let J, G be (multi-)graphs. A *model* $(\mathcal{U}, \mathcal{E})$ of J in G is a collection $\mathcal{U}$ of disjoint, connected sets $U_x \subseteq V(G), x \in V(J)$, and a collection $\mathcal{E}$ of internally disjoint U_x–U_y paths E_e, one for each edge $e = xy$ of J, such that E_e is disjoint from every U_z with $z \neq x, y$. The U_x are the *branch sets* and the E_e are the *branch paths* of the model. A model $(\mathcal{U}, \mathcal{E})$ of J in G is *K-fat* for $K \in \mathbb{N}$ if $\mathrm{dist}_G(Y, Z) \geq K$ for every two distinct $Y, Z \in \mathcal{U} \cup \mathcal{E}$ unless $Y = E_e$ and $Z = U_x$ for some vertex $x \in V(J)$ incident to $e \in E(J)$, or vice versa. We say that J is a *(K-fat) minor* of G, if G contains a (K-fat) model of X. We remark that the 0-fat minors of G are precisely its minors.

Lemma 2.1. *Let J, G be (multi-)graphs, and let $\dot{J}$ be the graph obtained from J by subdividing each of its edges precisely once. If J is a $3K$-fat minor of G for some $K \in \mathbb{N}$, then $\dot{J}$ is a K-fat minor of G.*

This lemma is a variant of [13, Lemma 5.3]; we include a proof for convenience.

Proof. Let $(\mathcal{U}, \mathcal{E})$ be a $3K$-fat model of J in G. We construct a K-fat model $(\mathcal{U}', \mathcal{E}')$ of $\dot{J}$ in G as follows. For every $x \in V(J)$, we keep $U'_x := U_x$ as a branch set. For every edge $e = xy \in E(J)$, we let u_e be the last vertex on E_e, as we move from U_x to U_y along E_e, such that $d_G(U_x, u_e) \leq K$, and we let v_e be the first vertex after u_e along E_e such that $d_G(U_y, v_e) \leq K$. We let the branch set U'_{w_e} for the subdivision vertex of $\dot{J}$ on e be the subpath of E_e between u_e and v_e. For $z \in \{x, y\}$, we let E'_{zw_e} be an U'_z–U'_{w_e} path of length K. This completes the definition of $(\mathcal{U}', \mathcal{E}')$.

As $(\mathcal{U}, \mathcal{E})$ is $3K$-fat and $E'_{xw_e} \subseteq B_G(E_e, K)$ for all edges of $\dot{J}$, we have $d_G(E'_{xw_e}, E'_{yw_f}) \geq 3K - 2K = K$ for all edges $xw_e \neq yw_f$ of $\dot{J}$, unless $e = f$, in which case we have $d_G(E'_{xw_e}, E'_{yw_e}) \geq d_G(U_x, U_y) - ||E'_{xw_e}|| - ||E'_{yw_e}|| = 3K - K - K = K$ by the choice of the branch paths of $\dot{J}$. Similarly and because $U'_{w_e} \subseteq E_e$ for all subdivision vertices of $\dot{J}$, we have $d_G(U'_x, U'_y) \geq 3K$ for all $x \neq y \in V(\dot{J})$, unless one of x, y is a subdivision vertex w_e on an edge e of J incident with the other, in which case we have $d_G(U'_x, U'_y) \geq K$ by the choice of the U'_{w_e}. Hence, it remains to consider $x \in V(\dot{J})$ and $yw_e \in E(\dot{J})$. If x is a subdivision vertex on an edge f of J, then $d_G(U'_x, E'_{yw_e}) \geq d_G(E_f, E_e) \geq 3K - K = 2K$. Otherwise, $d_G(U'_x, E'_{yw_e}) \geq d_G(U_x, U_y) - ||E'_{yw_e}|| = 3K - K = 2K$, as desired. □

2.3 Quasi-isometries and graph-partitions

Let G, H be graphs. For $M \in \mathbb{R}_{\geq 1}$ and $A \in \mathbb{R}_{\geq 0}$, an *$(M, A)$-quasi-isometry* from G to H is a map $\varphi : V(G) \to V(H)$ such that

(Q1) $M^{-1} \cdot d_G(u, v) - A \leq d_H(\varphi(u), \varphi(v)) \leq M \cdot d_G(u, v) + A$ for every $u, v \in V(G)$, and

(Q2) for every $h \in V(H)$ there is $v \in V(G)$ such that $d_H(h, \varphi(v)) \leq A$.

We say that a map $\varphi : V(G) \to V(H)$ has *multiplicative distortion* M (respectively, *additive distortion* A) if it satisfies (Q1) with $A = 0$ (resp. $M = 1$).

A *graph-partition of G over H*, or *H-partition* for short, is a partition $\mathcal{H} := (V_h : h \in V(H))$ of $V(G)$ indexed by the nodes of H such that for every edge $uv \in E(G)$, if $u \in V_g$ and $v \in V_h$, then $g = h$ or $gh \in E(H)$. (This notion generalizes tree-partitions.)

We say that $\mathcal{H}$ is *honest*, if V_h is non-empty for all $h \in V(H)$ and if for every edge $gh \in E(H)$ there exists an edge $uv \in V(G)$ such that $u \in V_g$ and $v \in V_h$. We say that $\mathcal{H}$ is *R-bounded*, if each V_h has diameter at most $R(K)$.

Lemma 2.2. *Let H, G be graphs, and let $\mathcal{H}$ be an honest, R-bounded H-partition of G for some $R \in \mathbb{R}$. Then G is $(R+1, R/(R+1))$-quasi-isometric to H.*

This is a special case of [3, Lemma 3.9]; we include a proof for convenience:

Proof. As the V_h are pairwise disjoint and cover $V(G)$, there is for every $v \in V(G)$ a unique $h_v \in V(H)$ such that $v \in V_{h_v}$. We claim that $\varphi : V(G) \to V(H)$ with $\varphi(v) := h_v$ is the desired quasi-isometry from G to H. Let us check that φ satisfies both properties of the definition of quasi-isometry:

(Q2): As the V_h are non-empty, there is for every $h \in V(H)$ some $v \in V(G)$ such that $h = \varphi(v)$, and hence h has distance 0 from $\varphi(v)$.

(Q1): Fix $u, v \in V(G)$. Since $w \in V_{\varphi(w)}$ for all $w \in V(H)$, every u–w path P in G of length $\ell \in \mathbb{N}$ induces a $\varphi(u)$–$\varphi(w)$ walk in H of length at most ℓ with vertex set $\{h \in V(H) \mid \exists p \in V(P) : p \in V_h\}$. Hence, $d_H(\varphi(u), \varphi(v)) \leq d_G(u, v)$.

Conversely, every $\varphi(u)$–$\varphi(v)$ path in H of length ℓ can be turned into a u–v walk in G of length at most $\ell \cdot (R+1) + R$ as the V_h have diameter at most R and $\mathcal{H}$ is honest. Hence, $d_G(u,v) \leq (R+1) \cdot d_H(\varphi(u), \varphi(v)) + R$. □

3 Structure of the proof of Theorem 1.2

For the proof of Theorem 1.2 we construct a graph-partition of a graph G with no K-fat $K_{2,t}$ minor, and then employ Lemma 2.2 to obtain the desired quasi-isometry. More precisely, we will prove the following stronger version of Theorem 1.2:

Theorem 3.1. *For every $t \in \mathbb{N}$ there exists a function $R : \mathbb{N} \to \mathbb{N}$ such that every graph G with no K-fat $K_{2,t}$ minor has an honest, $R(K)$-bounded graph-partition over a graph H such that every 2-connected multi-graph which is a minor of H is a K-fat minor of G.*

Let us first show that Theorem 3.1 implies Theorem 1.2:

Proof of Theorem 1.2 given Theorem 3.1. Fix $t, K \in \mathbb{N}$, and let G be a graph with no K-fat $K_{2,t}$ minor. Let $(H, (V_h)_{h \in V(H)})$ be an R-bounded graph-partition of G as provided by Theorem 3.1. Then G is $(R+1, R/(R+1))$-quasi-isometric to H by Lemma 2.2, and H has no $K_{2,t}$ minor. □

In this proof of Theorem 1.2 we showed that G is quasi-isometric to the graph H from Theorem 3.1. Since H has the property that all its 2-connected minors are K-fat minors of G, we have the following corollary:

Corollary 3.2. *Fix $t \in \mathbb{N}$, and let $\mathcal{J}$ be a class of finite, 2-connected graphs containing $K_{2,t}$. Then there exists a function $f : \mathbb{N} \to \mathbb{N}^2$ such that every graph with no K-fat minor in $\mathcal{J}$ is $f(K)$-quasi-isometric to a graph with no minor in $\mathcal{J}$.* □

Our proof of Theorem 3.1 will be divided into two steps. The first step is to structure our graph G as an H-partition as in Lemma 2.2, but with additional properties (Lemma 3.4 below). The second step is to show that these properties imply that any 2-connected subgraph of H is a K-fat minor of G (Lemma 3.3). To describe these additional properties ((i)–(iv) below), we need the following definitions.

A *rooted* graph is a pair (H, s) where H is a graph and s is one of its vertices, called its *root*. We will sometimes omit s from the notation if it is clear from the context. A rooted graph (H, s) has a natural layering: we denote by $L^i = L^i_{H,s} := \{h \in V(H) : d_H(s,h) = i\}$ the *i-th layer* of H. Given a vertex $h \in V(H)$ we denote by $i_h = i_{h,s}$ the unique integer satisfying $h \in L^{i_h}$.

Let $\mathcal{H} = (H, (V_h)_{h \in V(H)})$ be a graph-partition of a graph G over a graph H. If H is rooted, then for every $n \in \mathbb{N}$ we let $G^n = G^n_{\mathcal{H}}$ denote the subgraph of G induced by those vertices that are contained in partition classes V_h of nodes h in the layers of H up to L^n, i.e. $G^n := G[\bigcup_{i \leq n} \bigcup_{h \in L^i} V_h]$.

All graphs H used in graph-partitions $\mathcal{H} = (H, (V_h)_{h \in V(H)})$ in the remainder of this paper will be rooted, and we will ensure that

(i) for all $i \in \mathbb{N}$ the layer L^i is an independent set, i.e. there are no edges $xy \in E(H)$ with $x, y \in L^i$. In particular, H is bipartite, and for every edge $gh \in E(H)$ there exists $i \in \mathbb{N}$ such that $g \in L^i$ and $h \in L^{i+1}$.

Given $\mathcal{H}$ as above, and a node h of H, we let $\partial_h^{\downarrow}$ be the set of vertices of V_h that send an edge to some vertex of G^{i_h-1}.

The *height* R_h of a node h of H is the maximum distance $\max_{v \in V_h} d(\partial_h^{\downarrow}, v)$ of one of its vertices from its 'bottom' $\partial_h^{\downarrow}$. We say that V_h is *level*, if

(ii) $V_h = B_{G-G^{i_h-1}}(\partial_h^{\downarrow}, R_h)$.

Recall that we are trying to produce a graph-partition $\mathcal{H}$ of our graph G as in Theorem 3.1, so that every 2-connected minor J of H is a K-fat minor of G. The naive way to try to turn $J < H$ into a K-fat minor of G is to replace each vertex $h \in V(H)$ in the model of J by V_h. But this is too naive for two reasons: firstly, the V_h are not necessarily connected, and secondly, they are not necessarily K-far apart when we want them to be. To address these issues, instead of using a V_h in our branch sets, we will instead use a connected region of G around $\partial_h^{\downarrow}$. This region (depicted in (dark) blue in Figure 1) will consist of a subgraph of V_h of height less than $R_h - K$, as well as an *undergrowth*, i.e. a subgraph of the layer below i_h (hence outside V_h) used to ensure connectedness. We use the following notation to describe these subgraphs precisely. For $h \in V(H)$ and $R \in \mathbb{N}$, let

$$\partial_h^{\uparrow}(R) := B_{G-G^{i_h-1}}(\partial_h^{\downarrow}, R).$$

In particular, (ii) can be reformulated as $V_h = \partial_h^{\uparrow}(R_h)$, but we will use this notation with $R < R_h$ to capture a shorter subgraph of V_h. To define the aforementioned undergrowth, we similarly introduce

$$\partial_h^{\downarrow}(r) := B_G(\partial_h^{\downarrow}, r) \setminus \partial_h^{\uparrow}(r)$$

for $h \in V(H)$ and $r \in \mathbb{N}$. We remark that we think of $\partial_h^{\downarrow}(r)$ as lying 'below' $\partial_h^{\downarrow}$ and being mostly contained in G^{i_h-1}. In fact, whenever we use $\partial_h^{\downarrow}(r)$, we will make sure that for most other nodes $g \in V(H)$ in the same layer as h, their $\partial_g^{\downarrow}$ is more than r far apart from $\partial_h^{\downarrow}$, so that $\partial_h^{\downarrow}(r)$ cannot enter G^{i_h} through $\partial_g^{\downarrow}$ (and hence will be disjoint from V_g). (The only exception will be nodes $g \in L^{i_h}$ that can be separated from h by removing a single node of H (see (iv) below).)

The second step of our proof of Theorem 3.1 mentioned above is made precise by the following lemma (see Figure 1 for a sketch of the properties (ii) to (iv)):

Lemma 3.3. *Let $K, \ell \in \mathbb{N}$, let H be a rooted graph, and let G be a graph with an honest graph-partition $(H, (V_h)_{h \in V(H)})$ satisfying* (i) *and* (ii) *for every h. Suppose every $h \in V(H)$ has height $R_h \geq \ell + K$, and there is $r_h \in \mathbb{N}$ with $0 < r_h \leq \ell$ such that*

(iii) $\partial_h^{\uparrow}(R_h - \ell - K) \cup \partial_h^{\downarrow}(r_h)$ *is connected, and*

(iv) *for all non-adjacent $g \neq h \in V(H)$ either $d_G(V_g, V_h) \geq 2 \cdot \max\{r_g, r_h\} + 3K$, or there is a node in H that separates g, h.*

Then every 2-connected subgraph of H is a K-fat minor of G.

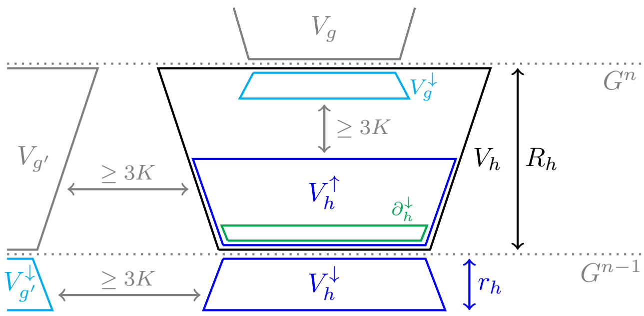

Figure 1: Depicted is a partition class V_h of the graph-partition in Lemma 3.3. The (dark) blue vertex set $V_h^{\uparrow} \cup V_h^{\downarrow}$ is connected by (iii), and $d_G(V_h \cup V_h^{\downarrow}, V_{g'} \cup V_{g'}^{\downarrow}) \geq 3K$ holds by (iv).

Given the setup of this lemma, let $V_h^{\uparrow} := \partial_h^{\uparrow}(R_h - \ell - K)$ and $V_h^{\downarrow} := \partial_h^{\downarrow}(r_h)$. Thus (iii) says that $V_h^{\uparrow} \cup V_h^{\downarrow}$ is connected.

Let us briefly sketch how Lemma 3.3 is proved. Given a 2-connected $J \subseteq H$, we build a model of J in G by replacing each vertex $h \in V(J)$ by $V_h^{\uparrow} \cup V_h^{\downarrow}$, which is connected by (iii) as just mentioned. For each edge $e = hg \in E(J)$ where g is in the layer above that of h, we model e by a branch path within V_h incident with the undergrowth $V_g^{\downarrow}$ of g inside V_h. We have tuned our parameters (by demanding $r_h \leq \ell$) so that each V_h has a *buffer zone* above $V_h^{\uparrow}$ and below all undergrowths protruding from the layer above, where it is safe to choose the branch paths (which are geodesics of length K). We then use (iv) to show that the branch sets in G are pairwise far apart.

The final step in the proof of Theorem 3.1 will then be to show that if a graph does not contain $K_{2,t}$ as a fat minor, then it has a graph-partition satisfying (i) to (iv) whose partition classes all have small radius. In fact, it will be more convenient to exclude Θ_t as a fat minor, where Θ_t denotes the multi-graph on two vertices with t parallel edges. Note that $K_{2,t}$ can be obtained from Θ_t by subdividing each of its edges precisely once.

Lemma 3.4. *There exists a function $R : \mathbb{N}^2 \to \mathbb{N}$ satisfying the following. Let $t, K \in \mathbb{N}$ with $t \geq 3$, and let G be a graph with no K-fat Θ_t minor. Then G admits an $R(t,K)$-bounded, honest graph-partition satisfying* (i) *to* (iv) *for some $\ell \in \mathbb{N}$.*

Together, Lemmas 3.3 and 3.4 imply Theorem 3.1:

Proof of Theorem 3.1. For $t = 0$, every graph excluding $K_{2,0}$ as a fat minor has bounded radius, and hence the assertion follows trivially. For $t = 1$, it is easy to see that every graph excluding $K_{2,1}$ as fat minor consists only of components that each have bounded diameter, and hence the assertion follows trivially. For $t = 2$, the result follows from (the proof of) the K_3 case of Conjecture 1.1 (see [13, Theorem 3.1]) and Lemma 2.1, where we note that in this case, G admits a tree-partition over a tree T, which has no 2-connected minors. Hence, we may assume $t \geq 3$.

Since $K_{2,t}$ is not a K-fat minor of G, it follows by Lemma 2.1 that Θ_t is not a $3K$-fat minor of G. Let $(H, (V_h)_{h \in V(H)})$ be the graph-partition provided by Lemma 3.4 for $G, t, 3K$. Let J be a 2-connected (multi)-graph that is a minor of H, and let J' be an $\subseteq$-minimal subgraph of H which still contains J as a minor. It is straight forward to check that J' is 2-connected. By Lemma 3.3, J' is a K-fat minor of G, and so J is a K-fat minor of G. □

4 Proof of Lemma 3.3

For every $h \in V(H)$, recall that

$$V_h^{\uparrow} := \partial_h^{\uparrow}(R_h - \ell - K), \text{ and}$$
$$V_h^{\downarrow} := \partial_h^{\downarrow}(r_h)$$

(see Figure 1). In particular, $V_h^{\uparrow} \cup V_h^{\downarrow}$ is connected by (iii). Let also $L^i := L^i_{H,s}$, for $i \in \mathbb{N}$, denote the i-th layer of H with respect to its root s.

Let J be a 2-connected subgraph of G. Our aim is to find a K-fat model of J in G, and we start with the branch paths. Let $f = gh \in E(J) \subseteq E(H)$. By (i), we may assume that $h \in L^i$ and $g \in L^{i+1}$ for some $i \in \mathbb{N}$. Since $\mathcal{H}$ is honest, there exists an edge $uv \in E(G)$ with $u \in V_h$ and $v \in V_g$, and hence $V_g^{\downarrow} \cap V_h \neq \emptyset$ since $r_g > 0$. We choose a $V_g^{\downarrow}$–$V_h^{\uparrow}$ path $Q^f = q_0^f \dots q_{\ell_f}^f$ through V_h of length $\ell_f := d_{G[V_h]}(V_g^{\downarrow}, V_h^{\uparrow})$, which exists by (ii) (Figure 1). Note that $\ell_f = \ell + K - r_h$ as $V_h^{\uparrow} = \partial_h^{\uparrow}(R_h - \ell - K)$ and $V_g^{\downarrow} = \partial_g^{\downarrow}(r_g)$ where $\partial_g^{\downarrow}$ is contained in the neighbourhood of G^{i_h}. In particular, $\ell_f \geq K$ since $r_h \leq \ell$. We declare the initial segment $E_f := q_0^f \dots q_K^f$ of length K of Q^f to be the branch path corresponding to f. The remaining subpath $T_f := q_K^f \dots q_{\ell_f}^f$ will be called the *tentacle* of f, and we will make it part of the branch set below, to ensure that each branch path attaches to the branch sets of its end-vertices (see Figure 2).

To complete our construction of a model of J in G, we now define the branch sets U_x as follows: for each $x \in V(J)$, let F_x be the set of edges of J that are incident with x and whose other endvertex lies in L^{i_x+1}, and let (see Figure 2)

$$U_x := V_x^{\uparrow} \cup V_x^{\downarrow} \cup \bigcup_{e \in F_x} T_e \subseteq V_x \cup V_x^{\downarrow}.$$

We claim that these U_x and E_e form the branch sets and branch paths of a K-fat model of J in G.

By (iii) and because $q_{\ell_e}^e, \in V_h^{\uparrow}$ for all $e \in F_x$, the sets U_x are connected. Since, by definition, every branch path E_e of an edge $e = xy \in E(J)$, with $i_x < i_y$, starts in $q_0^e \in V_y^{\downarrow} \subseteq U_y$ and ends in $q_K^e \in V(T_{xy}) \subseteq U_x$, it follows that $((U_x)_{x \in V(J)}, (E_e)_{e \in E(J)})$ is a model of J once we have shown that all pairs of non-incident branch sets and/or paths are disjoint. We will prove that they are even K-far apart in G, showing that our model of J is K-fat.

For this, let us first note that since J is 2-connected, it follows by (iv) that

$$d_G(V_x, V_y) \geq 2 \cdot \max\{r_x, r_y\} + 3K \qquad (*)$$

for all $x, y \in V(J) \subseteq V(H)$ with $xy \notin E(H)$. In particular,

$$d_G(V_x \cup V_x^{\downarrow}, V_y \cup V_y^{\downarrow}) \geq 3K \qquad (**)$$

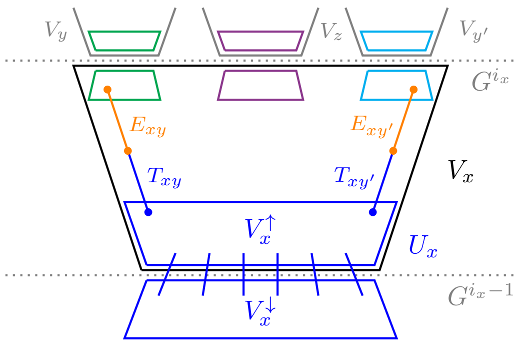

Figure 2: Depicted is an illustration of U_x, where $y, y', z \in V(J)$ and $xy, xy' \in E(J)$ and $xz \notin E(J)$.

for all $x, y \in V(J) \subseteq V(H)$ with $xy \notin E(H)$.

Let $e = xy, e' = x'y' \in E(J)$ be distinct edges of J. Since H is bipartite by (i) and hence triangle-free, and because $e \neq e'$, it follows that there are $a \in \{x, y\}$ and $b \in \{x', y'\}$ such that $a \neq b$ and $ab \notin E(H)$. Thus, by $(**)$ and because E_e and $E_{e'}$ meet $V_a \cup V_a^{\downarrow}$ and $V_b \cup V_b^{\downarrow}$, respectively, we have that

$$d_G(E_e, E_{e'}) \geq d_G(V_a \cup V_a^{\downarrow}, V_b \cup V_b^{\downarrow}) - ||E_e|| - ||E_{e'}|| \geq 3K - K - K = K.$$

Now let $z \in V(J)$ and $e = xy \in E(J)$ such that $z \notin \{x, y\}$. Once again, because H is triangle-free, there exists $a \in \{x, y\}$ such that $za \notin E(H)$. Hence, as above,

$$d_G(U_z, E_e) \geq d_G(V_z \cup V_z^{\downarrow}, E_e) \geq d_G(V_z \cup V_z^{\downarrow}, V_a \cup V_a^{\downarrow}) - ||E_e|| \geq 3K - K \geq K.$$

Finally, let $x \neq y \in V(J)$. If $xy \notin E(H)$, then, by $(**)$,

$$d_G(U_x, U_y) \geq d_G(V_x \cup V_x^{\downarrow}, V_y \cup V_y^{\downarrow}) \geq 3K \geq K,$$

where we used that $U_z \subseteq V_z \cup V_z^{\downarrow}$ for all $z \in V(J)$.

So we may assume that $xy \in E(H)$. Then by (i) and without loss of generality, $x \in L^{i-1}$ and $y \in L^i$ for some $i \in \mathbb{N}$. By (ii), it follows that

$$d_G(V_x^{\uparrow} \cup V_x^{\downarrow}, V_y \cup V_y^{\downarrow}) \geq d_G(V_x^{\uparrow}, \partial_y^{\downarrow}) - r_y \geq (\ell + K) - \ell = K.$$

It remains to show that $d_G(T_e, V_y \cup V_y^{\downarrow}) \geq K$ for all edges $e \in F_x$. (Recall that all tentacles of y are contained in V_y.) For this, let $e = xz \in E(J)$ with $e \in F_x$ be given. So $z \in L^i$. We split T^e into an 'upper part' $T_1^e := V(T^e) \cap B_G(\partial_z^{\downarrow}, r_y + K)$ and a 'lower part' $T_0^e := V(T^e) \setminus B_G(\partial_z^{\downarrow}, r_y + K)$. Note that T_1^e is empty if $r_z \geq r_y$, which is in particular the case when $y = z$. We show separately that both T_1^e, T_0^e have distance at least K from $V_y \cup V_y^{\downarrow}$. Indeed, if T_1^e is non-empty (and hence $z \neq y$), we have

$$d_G(T_1^e, V_y \cup V_y^{\downarrow}) \geq d_G(B_G(\partial_z^{\downarrow}, r_y + K), V_y \cup V_y^{\downarrow}) \geq d_G(V_z, V_y) - (r_y + K) - r_y$$

since $\partial_z^{\downarrow} \subseteq V_z$ and $V_y^{\downarrow} \subseteq B_G(V_y, r_y)$. Hence, by $(*)$,

$$d_G(T_1^e, V_y \cup V_y^{\downarrow}) \geq (2r_y + 3K) - r_y - K - r_y \geq K.$$

Moreover, by (ii) and since Q^e is a $V_z^{\downarrow}$–$V_x^{\uparrow}$ path of length $d_G(V_z^{\downarrow}, V_x^{\uparrow})$, we have $V(T_0^e) \subseteq B_G(V_x^{\uparrow}, R_x - r_y - K)$. It follows that

$$d_G(T_0^e, V_y \cup V_y^{\downarrow}) \geq d_G(V_x^{\uparrow}, V_y \cup V_y^{\downarrow}) - (R_x - r_y - K).$$

Since $V_y^{\downarrow} \subseteq B_G(V_y, r_y)$ and because $d_G(V_x^{\uparrow}, V_y) \geq R_x$ by (ii), it follows that

$$d_G(T_0^e, V_y \cup V_y^{\downarrow}) \geq d_G(V_x^{\uparrow}, V_y) - (R_x - r_y - K) - r_y = R_x - (R_x - K) = K.$$

This concludes the proof that $d_G(U_x, U_y) \geq K$ for all $x \neq y \in V(J)$, and hence completes the proof of Lemma 3.3. □

Corollary 4.1. *There is a polynomial-time algorithm that, given some $K \in \mathbb{N}$, a finite? graph G, an H-partition of G as in Lemma 3.3, and a 2-connected subgraph J of H, returns a K-fat model of J in G.*

Proof. The above proof is constructive, and provides an efficient procedure to turn a subgraph J of H into a K-fat model of J in H. □

5 Component structure and K-fat Θ_t minors

The rest of the paper is devoted to the proof of Lemma 3.4, for which we will construct a graph-partition of our graph G recursively. At the beginning of the n-th step of the recursion, we will already have constructed a graph-partition $\mathcal{H}^{n-1}$ of some induced subgraph G^{n-1} of G. To proceed with the construction, we need that the components C of $G - G^{n-1}$ satisfy two conditions. First, their boundaries $\partial_G C$ should not be too large, so that we can partition them into few sets of bounded radius. For this, we establish Lemma 5.2 below, which finds a fat Θ_t minor otherwise. Furthermore, we need that not too many components attach to the same bags of $\mathcal{H}^{n-1}$. For this, we establish Lemma 5.3 below, which again finds a fat Θ_t minor otherwise.

We start with a simpler lemma needed for both aforementioned lemmas.

Lemma 5.1. *Let G be a graph, and $K \in \mathbb{N}$. Let $X, Y \subseteq V(G)$ be connected and $d_G(X, Y) \geq K$. For every $t \in \mathbb{N}_{\geq 1}$, if $B_G(X, K) \cap V(Y)$ contains t vertices which are pairwise at least $3K$ apart, then Θ_t is a K-fat minor of G.*

Moreover, if G is finite, then there is a polynomial-time algorithm (for fixed t) that given the above data either confirms that no such t-tuple of vertices exists, or returns a K-fat Θ_t minor of G.

Proof. Assume that $B_G(X, K) \cap Y$ contains vertices $u_1, \ldots, u_t$ which are pairwise at least $3K$ apart in G. For every $i \in [t]$, let P_i be a u_i–X path of length K. Then $V_1 := Y$ and $V_2 := X$ form the branch sets and the P_i form the branch paths of a K-fat model of Θ_t in G. Indeed, we have $d_G(V_1, V_2) = d_G(X, Y) \geq K$ by assumption, and $d_G(P_i, P_j) \geq d_G(u_i, u_j) - ||P_i|| - ||P_j|| \geq 3K - K - K = K$.

For the second claim, it is straightforward to efficiently check if $B_G(X, K) \cap Y$ contains such a t-tuple, as there are at most n^t tuples to consider. If such a t-tuple is found, then the above proof provides an efficient procedure for finding a K-fat Θ_t minor. □

Lemma 5.2. *Let G be a graph, and let $X \subseteq V(G)$ be connected. Let further $K \in \mathbb{N}$, and let C be a component of $G - B_G(X, K-1)$. If Θ_t is not a K-fat minor of G for some $t \geq 2$, then $\partial_G C$ has at most $t-1$ $3K$-near-components and each of them has diameter less than $6K(t-1)$.*

Moreover, if G is finite, then there is a polynomial-time algorithm that either confirms that C has the aforementioned properties, or returns a K-fat Θ_t minor of G.

Proof. If $\partial_G C$ has at least t $3K$-near components, then taking one vertex from each $3K$-near component yields t vertices in $\partial_G C$ which are pairwise at least $3K$ apart. Applying Lemma 5.1 (with $X := X$ and $Y := V(C)$) yields that Θ_t is a K-fat minor of G.

Now suppose that some $3K$-near component C' of $\partial_G C$ has diameter at least $6K(t-1)$, and pick vertices $u, v \in V(C')$ with $d_G(u,v) \geq 6K(t-1)$. Since C' is a $3K$-near component, there exists a $3K$-near path $P = x_0 \dots x_n$ in C' from $u = x_0$ to $v = x_n$. Let W be an u–v walk in G obtained from P by adding for every $i \in \{0, \dots, n-1\}$ an x_i–x_{i+1} path of length at most $3K$ to P. Since $d_G(u,v) \geq 6K(t-1)$, the walk W has vertices $u = y_1, y_2, \dots, y_{t-1}, y_t = v$ which are pairwise at least $6K$ apart in G. By the definition of W, there exists for every y_j some x_{i_j} in P, which hence lies in $\partial_G C$, that has distance at most $3K/2$ from y_i. It follows that $d_G(x_{i_j}, x_{i_\ell}) \geq d_G(y_j, y_\ell) - d_G(y_j, x_{i_j}) - d_G(y_\ell, x_{i_\ell}) \geq 6K - 3K = 3K$. Thus, applying Lemma 5.1 (with $X := X$ and $Y := V(C)$) to the x_{i_j} for $j \in [t]$ yields that Θ_t is a K-fat minor of G.

For the second statement, it is again straightforward to compute and count the $3K$-near-components of $\partial_G C$, and to calculate their diameters, and so we can efficiently check whether C satisfies the desired properties. If not, and the number of these $3K$-near-components is at least t, then invoking Lemma 5.1 as above will return a K-fat Θ_t minor. Finally, if one of these $3K$-near-components C' has diameter at least $6K(t-1)$, then the above proof yields an efficient procedure for finding a t-tuple of vertices in C' pairwise at distance at least $3K$, and invoking Lemma 5.1 again returns a K-fat Θ_t minor. □

Another consequence of Lemma 5.1 is

Lemma 5.3. *Let $K, t, n \in \mathbb{N}$ with $t \geq 3$ and $n \leq t-1$, and let G be a graph with no K-fat Θ_t minor. Let $X_1, X_2, \dots, X_n$ be connected subsets of $V(G)$ that are pairwise at least $3K$ apart and set $V' := \bigcup_{i \in [n]} B_G(X_i, K-1)$. Let $\mathcal{C}$ be the set of components of $G - V'$ that each have neighbours in at least two distinct $B_G(X_i, K-1)$. Then there is no set of more than $(t-1)^3(t-2)$ vertices of $\bigcup_{C \in \mathcal{C}} \partial_G C$ pairwise at distance at least $3K$.*

Moreover, if G is finite, then there is a polynomial-time algorithm that either confirms that $\mathcal{C}$ has the aforementioned property, or returns a K-fat Θ_t minor of G.

Proof. Suppose for a contradiction that there is a set $U \subseteq \bigcup_{C \in \mathcal{C}} \partial_G C$ of size at least $(t-1)^3(t-2)+1$ such that $d_G(u,u') \geq 3K$ for all $u, u' \in U$. For every $u \in U$, let $C_u \in \mathcal{C}$ be the component of $G - V'$ containing u.

By the pigeonhole principle and because $n \leq t-1$, there is $i \in [n]$ and a subset $U' \subseteq U$ of size at least $(t-1)^2(t-2)+1$ such that every $u \in U'$ has a neighbour in $B_G(X_i, K-1)$. Further, by the same argument and because every $C_u \in \mathcal{C}$ has neighbours in at least two distinct $B_G(X_j, K-1)$, it follows that there is $j \neq i \in [n]$ and a set $U'' \subseteq U'$ of size at least $(t-1)^2+1$ such that for every $u \in U''$ the component C_u has a neighbour in $B_G(X_j, K-1)$. Moreover, by Lemma 5.1 (applied to $X := X_i$ and $Y := V(C_u)$ for every $u \in U''$) and because Θ_t is not a K-fat minor of G, we deduce that there is a subset $W \subseteq U''$ of size at least t such that $C_u \neq C_{u'}$ for all $u \neq u' \in W$.

We now use W to show that Θ_t is a K-fat minor of G, which contradicts our assumptions and thus concludes the proof. For every $u \in W$ pick a u–X_i path Q_u of length K, which exists since $u \in N_G(B_G(X_i, K-1))$. Then by the choice of W, the paths Q_u form the branch paths and $V_1 := X_i$ and $V_2 := B_G(X_j, K-1) \cup \bigcup_{u \in W} V(C_u)$ form the branch sets of a model of Θ_t (see Figure 3). We claim that this model is K-fat. Indeed, we have

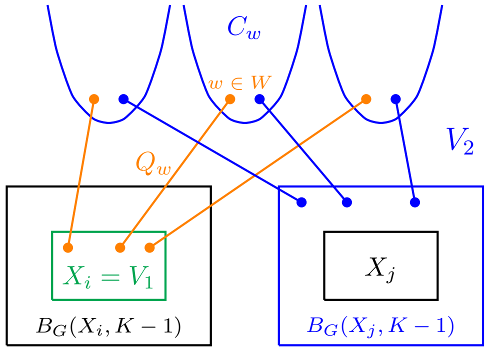

Figure 3: An illustration of the fat Θ_t minor in the proof of Lemma 5.3. The green and blue sets are its branch sets, and the orange paths are its branch paths.

$$d_G(Q_u, Q_{u'}) \geq d_G(u, u') - ||Q_u|| - ||Q_{u'}|| \geq 3K - K - K = K,$$

since $u, u' \in U$ and hence $d_G(u, u') \geq 3K$ by the assumption on U. Moreover,

$$d_G(V_1, B_G(X_j, K)) \geq d_G(X_i, X_j) - K \geq 3K - K > K$$

by the assumption on the X_k. Finally, we have $d_G(V_1, C_u) \geq K$ for all $u \in W$ since C_u is a component of $G - V'$, which concludes the proof.

For the second statement, it is straightforward to compute the set $\mathcal{C}$ of components of $G - V'$ that each have neighbours in at least two distinct sets $B_G(X_i, K-1)$, and to check if $\bigcup_{C \in \mathcal{C}} \partial_G C$ has such a t'-tuple of vertices where $t' := (t-1)^3(t-2)+1$, as there are at most $n^{t'}$ tuples to consider. If such a t'-tuple is found, then the above proof provides an efficient procedure for finding a K-fat Θ_t minor (where it might find the K-fat Θ_t minor by invoking Lemma 5.1, which is the only point in the proof (except for the contradiction in the end) where we used the assumption that Θ_t is not a K-fat minor of G). $\square$

6 A merging lemma

Recall that for the proof of Lemma 3.4 we will construct a graph-partition of a graph G recursively. After each step, we will have constructed a graph-partition $\mathcal{H}^n$ of some subgraph of G. In the next step, we will consider, for some suitable $K' \in \mathbb{N}$, the K'-near-components of the boundaries $\partial_G C$ of the remaining components C as candidates for the new partition classes which we aim to add to $\mathcal{H}^n$. However, some of the near-components might be too close to each other for (iv), in which case we combine them into one new partition class. The following lemma formalises this merging procedure.

Given a set U and partitions $\mathcal{P}, \mathcal{Q}$ of U, we say that $\mathcal{P}$ is a *coarsening* of $\mathcal{Q}$ if every $B \in \mathcal{Q}$ is a subset of some $A \in \mathcal{P}$.

Lemma 6.1. *Let $n \in \mathbb{N}$, let G be a graph, and let $\mathcal{Q}$ be a set of at most n disjoint subsets of $V(G)$. (We think of $\mathcal{Q}$ as a partition of $\bigcup \mathcal{Q}$.) Then for every $d, r \in \mathbb{N}$, there exist some $L \in \mathbb{N}$ with $r \leq L \leq r + \lfloor \frac{nd}{2} \rfloor$, and a coarsening $\mathcal{P}$ of $\mathcal{Q}$ such that*

(i) *for every $A \in \mathcal{P}$ and every $u, v \in A$ there is a sequence $(B_i)_{i \in [k]} \subseteq \mathcal{Q}$ with $B_i \subseteq A$ for all $i \in [k]$ such that $u \in B_1$, $v \in B_k$ and $d_G(B_{i-1}, B_i) \leq 2L$ for all $i \in \{2, \ldots, k\}$,*

(ii) *$d_G(A, A') \geq 2L + d$ for all $A \neq A' \in \mathcal{P}$, and*

(iii) *if $\operatorname{diam}_G(B) \leq D$ for all $B \in \mathcal{Q}$ and some $D \in \mathbb{N}$, then every $A \in \mathcal{P}$ has diameter at most $nD + (n-1)(2r + nd)$.*

Moreover, if G is finite, then $\mathcal{P}$ can be computed in polynomial time.

Proof. We first construct a coarsening $\mathcal{P}$ satisfying (i) and (ii), and then verify that $\mathcal{P}$ also satisfies (iii). We construct $\mathcal{P}$ recursively as follows. Set $\mathcal{P}_0 := \mathcal{Q}$ and $L_0 := r$, and assume that we have already defined $\mathcal{P}_m$ for some $m < n$ such that $\mathcal{P}_m$ has $n - m$ elements and satisfies (i) with $L_m := r + \lfloor \frac{md}{2} \rfloor$ instead of L. If $\mathcal{P}_m$ also satisfies (ii) with L_m instead of L, then $\mathcal{P} := \mathcal{P}_m$ and $L := L_m$ are as desired. In particular, if $m = n - 1$, then $|\mathcal{P}_m| = 1$, and hence $\mathcal{P}_m$ satisfies (ii) trivially.

Otherwise, pick two sets $A, A' \in \mathcal{P}_m$ with $d_G(A', A') < 2L_m + d$. Then $\mathcal{P}_{m+1} := (\mathcal{P}_m \setminus \{A, A'\}) \cup \{A \cup A'\}$ has $n - m - 1$ elements, and it still satisfies (i) with $L_{m+1} := r + \lfloor \frac{(m+1)d}{2} \rfloor \geq L_m + \lfloor \frac{d}{2} \rfloor$ instead of L. Indeed, let $a \in A$ and $a' \in A'$ such that $d_G(a, a') < 2L_m + d$. Then for every $u \in A$ and $v \in A'$ we can concatenate the sequences given by (i) for $u, a \in A$ and $a', v \in A'$, which yields a sequence for $u, v \in A \cup A'$ as in (i). This completes the construction of $\mathcal{P}$ and the verification that $\mathcal{P}$ satisfies (i) and (ii).

To check (iii), let $A \in \mathcal{P}$, and assume that $\operatorname{diam}_G(B) \leq D$ for some $D \in \mathbb{N}$ and all $B \in \mathcal{Q}$. Then

$$\operatorname{diam}_G(A) \leq nD + (n-1)2L.$$

by picking $u, v \in A$, and a sequence of $B_i's$ as in (i), and noting that we have at most n such $B_i's$. The right hand side is at most $nD + (n-1)(2r + nd)$ by our bound on L.

Since this recursive construction terminates after at most n steps, each of which only compares distances between pairs of at most n sets of vertices, it can be carried out by a polynomial-time algorithm. □

7 Proof of Lemma 3.4

We can now prove Lemma 3.4. For every $t, K \in \mathbb{N}$ with $t \geq 3$ set[2]

$$\begin{aligned}
N(t) &:= \left\lceil \frac{1}{2}(t-1)^3 \cdot (t-2) \right\rceil, \\
L(t,K) &:= \lceil 3K/2 \rceil + N(t) \cdot 3K, \\
L'(t,K) &:= N(t) \cdot \big(4 \cdot L(t,K) + 5K\big) + 2 \cdot L(t,K) + 3K, \\
R_0(t,K) &:= 3t^{12}K + 43t^9 K, \text{ and} \\
R(t,K) &:= R_0(t,K) + 2L'(t,K) \in O(t^{12}K).
\end{aligned}$$

We prove Lemma 3.4 with the function $R(t,K)$ and $\ell := L(t,K)$.

Let $t, K \in \mathbb{N}$, and let G be a graph with no K-fat Θ_t minor. By considering each component of G individually, we may assume that G is connected.

We construct the desired graph-partition $\mathcal{H} = (H, (V_h)_{h \in V(H)})$ of G recursively 'layer by layer', i.e. the nodes that we add to H in the n-th step of the construction will form the n-th layer $L^n := L^n_{H,s}$ of H with respect to the root s of H, which we specify in the first construction step.

Pick $o \in V(G)$ arbitrarily. We initialize $H^0 := (\{s\}, \emptyset)$ on a single vertex s, its root, and set $V_s := B_G(o, L'(t,K))$. Then $\mathcal{H}^0 = (H^0, (V_s))$ is an honest graph-partition of $G^0 = G[V_s]$. Moreover, $L^0 = \{s\}$.

Having defined graph-partitions $\mathcal{H}^i$ of G^i for every $i \leq n$, we proceed to construct $\mathcal{H}^{n+1}$. The main effort will go into finding a suitable partition $\mathcal{P}$ of $N_G(G^n)$ into sets of diameter at most $R_0(t,K)$ (whose construction we postpone for later). The new vertices of $H^{n+1} - H^n$ will be in bijection with the elements of $\mathcal{P}$. For each $A \in \mathcal{P}$, we introduce a vertex h_A, fix a 'height' $R_A = R_{h_A} \leq L'(t,K)$, and let $V_{h_A} := B_{G-G_n}(A, R_A)$ (thus ensuring that V_{h_A} is level, i.e. (ii) holds). We choose $\mathcal{P}$ and the heights R_A so that the V_{h_A} are pairwise disjoint, and there is no edge of G between V_{h_A} and V_{h_B} for $A \neq B \in \mathcal{P}$ (in fact, the V_{h_A} will be pairwise far apart; see (2) below). We add an edge between nodes $h, h' \in V(H^{n+1})$ whenever there is an edge in G between V_h and $V_{h'}$. By the last property, $L^{n+1} = V(H^{n+1} - H^n)$ is independent. Moreover, $L^n = V(H^n - H^{n-1})$ separates L^{n+1} from all L^i with $i \leq n-1$ since the partition classes of nodes $h \in L^n$ contain the neighbourhood of G^{n-1}. Hence, $V(H^{n+1} - H^n)$ is indeed the $(n+1)$st layer L^{n+1} of H^{n+1}. By definition, $\mathcal{H}^{n+1} = (H^{n+1}, (V_h)_{h \in V(H^{n+1})})$ is an honest graph-partition of $G^{n+1} = G[\bigcup_{h \in V(H^{n+1})} V_h]$.

If $N_G(G^n)$ is empty at some step n, which happens precisely when G has finite diameter, then the process terminates. This is the only difference between the finite and infinite diameter case throughout our proof.

We let $H := \bigcup_{n \in \mathbb{N}} H^n$. Then $\mathcal{H} := (H, (V_h)_{h \in V(H)})$ is an honest graph-partition of $\bigcup_{n \in \mathbb{N}} G^n$, which is equal to G since G is connected and each G^{n+1} contains the neighbourhood of G^n. By the comment above, $\mathcal{H}$ satisfies (i). Furthermore, $\mathcal{H}$ satisfies (ii) by the definition of V_{h_A} and because $\partial^{\downarrow}_{h_A} = A$. Moreover, as every $A \in \mathcal{P}$ has diameter at most $R_0(t,K)$ and $R_A \leq L'(t,K)$, every partition class V_{h_A} of $\mathcal{H}$ has diameter at most $R_0(t,K) + 2L'(t,K) = R(t,K)$, and hence $\mathcal{H}$ is $R(t,K)$-bounded.

[2] We remark that we rounded the function $R_0(t,K)$ up to make it more readable.

Thus, it only remains to specify $\mathcal{P}$ and the heights R_A, and check that (iii) and (iv) hold. We repeat these properties here: for all $h \in V(H)$

(1) $\partial_h^{\uparrow}(R_h - \ell - K) \cup \partial_h^{\downarrow}(r_h)$ is connected.

(We will specify the 'depths' r_h later on.) Recall that $\partial_h^{\downarrow}$ is the set of all vertices of V_h that have a neighbour in G^{i_h-1}; in particular, $\partial_{h_A}^{\downarrow} = A$ for every node $h_A \in L^{n+1}$ by definition. We need the following modified version of (iv):

(2) for all non-adjacent $g \neq h \in V(H)$ either $d_G(V_g, V_h) \geq 2 \cdot \max\{r_g, r_h\} + 3K$ or there is a node $x \in V(H)$ such that V_x separates V_g, V_h in G.

(Note that (2) immediately implies (iv) since $\mathcal{H}$ is honest.)

For our construction we need to inductively ensure that (1) and (2) hold for all $g, h \in V(H^n)$ and that the following property is true:

(3) Every component C of $G - G^{n-1}$ meets at most $t-1$ partition classes V_h of $\mathcal{H}^n$.

Letting $R_s := L'(t, K)$, $r_s := 0$, and $G^{-1} := \emptyset$ clearly satisfies (1) to (3) for $n = 0$.

For every component Z of $G - G^{n-1}$ let $\mathcal{D}_Z$ be the set of all components of $G - G^n$ that are contained in Z and that have neighbours in at least two distinct partition classes of $\mathcal{H}^n$. Recall that $\mathcal{C}(G - G^n)$ is the set of components of $G - G^n$. Let $\mathfrak{R}$ be the partition of $\mathcal{C}(G - G^n)$ comprising the $\mathcal{D}_Z$ as above and a singleton $\{C\}$ for each component C of $G - G^n$ that is not in any $\mathcal{D}_Z$ (i.e. that has neighbours in exactly one partition class of $\mathcal{H}^n$) (see Figure 4).

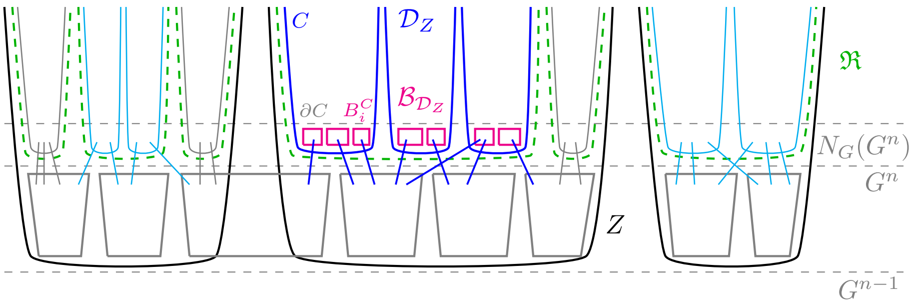

Figure 4: A visualisation of the partition $\mathfrak{R}$ of $\mathcal{C}(G-G^n)$ (in green, with dashed lines). Every partition class in $\mathfrak{R}$ is either a singleton comprising a component that has only neighbours in exactly one partition class of $\mathcal{H}^n$ (indicated in grey), or it is $\mathcal{D}_Z$ for some component Z of $G - G^{n-1}$ (indicated in light/dark blue).

Note that $\mathfrak{R}$ naturally induces a partition $\mathcal{R}$ of $N_G(G^n)$, by letting $\mathcal{R} := \{\partial_G(\bigcup \mathcal{D}) \mid \mathcal{D} \in \mathfrak{R}\}$. We will obtain $\mathcal{P}$ by refining $\mathcal{R}$.

For every $C \in \mathcal{C}(G - G^n)$, let $B_1^C, \ldots, B_{m_C}^C$ be the $3K$-near components of $\partial_G C$. We group these $3K$-near components together over $\mathcal{R}$ by considering

$$\mathcal{B}_{\mathcal{D}} := \{B_i^C : C \in \mathcal{D},\, i \leq m_C\} \text{ for every } \mathcal{D} \in \mathfrak{R}.$$

Set $\mathcal{B} := \bigcup_{\mathcal{D} \in \mathfrak{R}} \mathcal{B}_{\mathcal{D}}$, and note that $\bigcup \mathcal{B} = N_G(G^n)$. Our final partition $\mathcal{P}$ of $N_G(G^n)$ will be a refinement of $\mathcal{R}$ and a coarsening of $\mathcal{B}$.

We may think of $\mathcal{B}$ as candidate for the partition $\mathcal{P}$ of $N_G(G^n)$, and the B_i^C as candidates for the new partition classes V_h that we want to add to $\mathcal{H}^n$. By taking $r_h := 0$ for all such new V_h (and $R_h = 0$), they would already satisfy (2). However, we need that they also satisfy (1). But since the B_i^C are only $3K$-near components, they need not be connected. To make them connected, we might have to increase the heights R_h and 'depths' r_h to $3K/2$. By doing so, the B_i^C might no longer satisfy (2), in which case we have to merge B_i^C's that are to close. For this, we will employ Lemma 6.1, which ensures that the merged sets are far apart and have bounded diameters (see Figure 5). In order

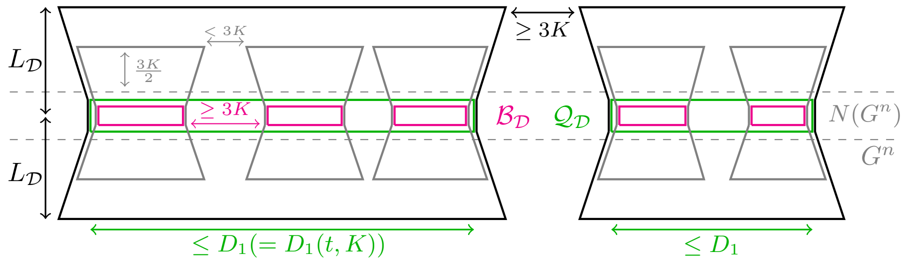

Figure 5: Indicated in pink is the partition $\mathcal{B}_\mathcal{D}$, whose partition classes B_i^C are pairwise at least $3K$ far apart. The grey boxes around the B_i^C are connected, but they no longer need to be pairwise $3K$ far apart. Applying Lemma 6.1 yields a coarsening $\mathcal{Q}_\mathcal{D}$ of $\mathcal{B}_\mathcal{D}$ such that the black boxes (of height $L_\mathcal{D}$) around the (green) partition classes of $\mathcal{Q}_\mathcal{D}$ are connected and pairwise at least $3K$ far apart. Moreover, the partition classes in $\mathcal{Q}_\mathcal{D}$ have bounded diameter.

to apply Lemma 6.1, we need to ensure that each $\mathcal{B}_\mathcal{D}$ contains only boundedly many elements all of bounded diameter. More precisely, we claim that for all $\mathcal{D} \in \mathfrak{R}$

$$\mathcal{B}_\mathcal{D} \text{ contains at most } (t-1)^3(t-2) \text{ elements,} \tag{$*$}$$

(this bound will play the role of n in our application of Lemma 6.1), and

$$\operatorname{diam}_G(B) \leq 6K(t-1) \text{ for every } B \in \mathcal{B}_\mathcal{D}. \tag{$**$}$$

For this, let $C \in \mathcal{C}(G - G^n)$. Then applying Lemma 5.2 to C (with some X that we specify in the next sentence) yields that every $3K$-near component B_i^C of $\partial_G C$ has diameter at most $6K(t-1)$ in G and that $m_C \leq t-1$; in particular, ($**$) holds. For this, let X be the component of $Y := G^n - B_G(G - G^n, K-1)$ which contains G^0. (Recall that $G^0 = G[V_s] = G[B_G(o, L'(t,K))]$, and hence G^0 is connected.) For the application of Lemma 5.2, we need to check that C is a component of $B_G(X, K-1)$, which we do next. Since C is connected and avoids X, it suffices to show that $N_G(C) \subseteq B_G(X, K-1)$. Pick $v \in N_G(C)$, and let $w \in Y$ be a vertex with distance precisely $K-1$ from v. We need to show that $w \in X$. Since $Y \subseteq G^n$, there are $i \leq n$ and $h \in L^i$ such that $w \in V_h$. Thus $w \in B_{G-G^{i-1}}(\partial_h^\downarrow, R_h)$ by the definition of V_h, and hence there is an v–$\partial_h^\downarrow$ path P in $G[V_h]$ of length $d_{G-G^{i-1}}(v, \partial_h^\downarrow)$. In particular, P is contained in Y as P is a subpath of an v–$\partial_h^\downarrow$ path of length $d_{G-G^{i-1}}(v, \partial_h^\downarrow) + K - 1$, of which only the first $K-1$ vertices (those not in P) are not contained in Y. Since all vertices in $\partial_h^\downarrow$ send an edge to G^{i-1} by the definition of $\partial_h^\downarrow$, there is a path in Y

from w to G^{i-1} (note that $V(G^{i-1}) \subseteq Y$). Inductively applying this argument thus yields that there is a w–G^0 path in Y, and thus $w \in X$ as claimed.

To complete the proof of $(*)$, note that if $\mathcal{D} = \{C\}$ for some $C \in \mathcal{C}(G - G^n)$, then $(*)$ follows immediately from $m_C \leq t - 1$. Otherwise, $\mathcal{D} = \mathcal{D}_Z$ for some component $Z \in \mathcal{C}(G - G^{n-1})$. We then obtain $(*)$ by applying Lemma 5.3 with the sets X_i being the sets $\partial_h^{\uparrow}(R_h - K - 1) \cup \partial_h^{\downarrow}(r_h)$ for nodes $h \in L^n$ whose partition class V_h has a neighbour in some $C \in \mathcal{D}$ (which implies that $\mathcal{D} = \mathcal{D}_Z$ is a subset of the set $\mathcal{C}$ from Lemma 5.3). For this, note that there are at most $t - 1$ such V_h by (3) and because the components in $\mathcal{D} = \mathcal{D}_Z$ are all contained in D. Moreover, note that the X_i are connected by (1) and are pairwise at least $3K$ apart by (2).

Having established the conditions $(*)$ and $(**)$, we can now apply Lemma 6.1 to each $\mathcal{B}_{\mathcal{D}}, \mathcal{D} \in \mathfrak{R}$, with the parameters being $n := |\mathcal{B}_{\mathcal{D}}| \leq 2N(t)$, $r := \lceil 3K/2 \rceil$, $d := 3K$ and $D := 6K(t-1)$. This merging yields a coarsening $\mathcal{Q}_{\mathcal{D}}$ of $\mathcal{B}_{\mathcal{D}}$ and some $L_{\mathcal{D}} \leq \ell$ (see Figure 5) such that every $A \in \mathcal{Q}_{\mathcal{D}}$ has diameter at most $D_1 := nD + (n-1)(2r + nd)$ (by (iii)) and such that $B_G(A, L_{\mathcal{D}})$ is connected (by (i) and because $B_G(B, \lceil 3K/2 \rceil)$ is connected for every $B \in \mathcal{B}_{\mathcal{D}}$ and because $L_{\mathcal{D}} \geq r = \lceil 3K/2 \rceil$). Moreover, (by (ii)) for all $A, A' \in \mathcal{Q}_{\mathcal{D}}$

$$d_G(A, A') \geq 2L_{\mathcal{D}} + 3K. \tag{$\Box$}$$

Set $\mathcal{Q} := \bigcup_{\mathcal{D} \in \mathfrak{R}} \mathcal{Q}_{\mathcal{D}}$, and note that $\bigcup \mathcal{Q} = \bigcup \mathcal{B} = N_G(G^n)$.

The partition $\mathcal{Q}$ is our new candidate for $\mathcal{P}$, and the $L_{\mathcal{D}}$ are our candidates for the 'heights' R_A. They would satisfy (3) and a variant of (2) (see $(\Box)$), and they would satisfy a variant of (1) with 'depths' $r_h := L_{\mathcal{D}}$ whereby we need the whole height for connectedness, i.e. for all $A \in \mathcal{Q}_{\mathcal{D}}$ we have that

$$B_G(A, L_{\mathcal{D}}) \text{ is connected,} \tag{$1'$}$$

which we have proven above. Note that $B_G(A, L_{\mathcal{D}})$ would be equivalent to $\partial_h^{\uparrow}(R_h) \cup \partial_h^{\downarrow}(r_h)$ if we would set $R_h, r_h := L_{\mathcal{D}}$ and $V_h := B_{G-G^n}(A, R_h)$. To achieve (1), we need to add a 'buffer zone' of height $\ell + K$, that is, we need to increase the 'height' R_A for each $A \in \mathcal{Q}$ by $\ell + K$. This increase in height might however violate (2) even if (2) was satisfied earlier, and therefore we need to perform another round of merging, namely to merge any sets in some $\mathcal{Q}_{\mathcal{D}}$ that violate (2), i.e. which are two close together (see Figure 6). This merging will ensure (2), and (1) will follow from $(1')$, as we will see below.

To perform the aforementioned merging, we now apply Lemma 6.1 again, to each $\mathcal{Q}_{\mathcal{D}}$ with $\mathcal{D} \in \mathfrak{R}$. More precisely, we apply Lemma 6.1 to $\mathcal{Q}_{\mathcal{D}}$ in the subgraph $G - G^n$ with $n' := |\mathcal{Q}_{\mathcal{D}}| \leq 2N(t)$, $r' := \ell + 2K$, $d' := 4\ell + 5K$ and $D' := D_1$. This yields a coarsening $\mathcal{P}_{\mathcal{D}}$ of $\mathcal{Q}_{\mathcal{D}}$ and some $L'_{\mathcal{D}} \leq L'(t, K) - \ell - K$ with $L'_{\mathcal{D}} \geq r'$ (see Figure 6). This new $L'_{\mathcal{D}}$ is the 'height' that we need to ensure connectedness as in $(1')$ (or in (1)), i.e. for all $A \in \mathcal{P}_d$ it follows by $(1')$ and Lemma 6.1 (i) that

$$B_{G-G^n}(A, L'_{\mathcal{D}}) \cup B_G(A, L_{\mathcal{D}}) \text{ is connected.} \tag{$1''$}$$

Moreover, by Lemma 6.1 (ii), for every $A \neq B \in \mathcal{P}_{\mathcal{D}}$,

$$d_{G-G^n}(A, A') \geq 2L'_{\mathcal{D}} + 4\ell + 5K. \tag{$\triangle$}$$

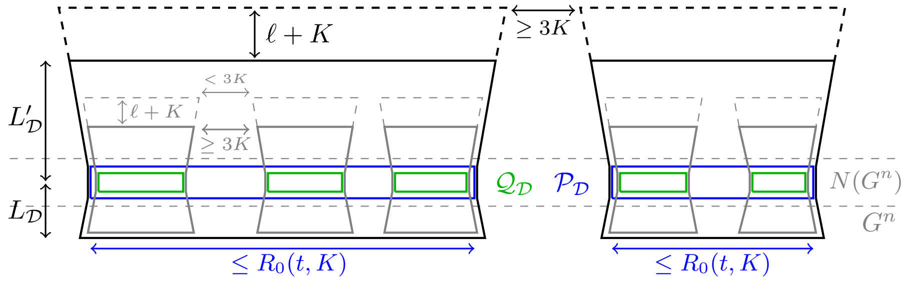

Figure 6: Indicated in green is the partition $\mathcal{Q}_{\mathcal{D}}$. The grey boxes around its partition classes (of height $L_{\mathcal{D}}$) are connected and pairwise at least $3K$ far apart, but to ensure (1), we need to add a 'buffer zone' of height $\ell + K$ (indicated with dashed lines). These taller boxes need no longer be pairwise $3K$ far apart. Applying Lemma 6.1 yields a coarsening $\mathcal{P}_{\mathcal{D}}$ of $\mathcal{Q}_{\mathcal{D}}$ such that the black boxes (of 'depth' $L_{\mathcal{D}}$ and height $L'_{\mathcal{D}}$) around the (blue) partition classes of $\mathcal{P}_{\mathcal{D}}$ are connected, and such that they are still $3K$ far apart even after adding a 'buffer zone' of height $\ell + K$. Moreover, the partition classes in $\mathcal{Q}_{\mathcal{D}}$ have diameter at most $R_0(t, K)$.

Setting $\mathcal{P} := \bigcup_{\mathcal{D} \in \mathfrak{R}} \mathcal{P}_{\mathcal{D}}$, we have defined our desired partition of $N_G(G^n)$. Note that $\mathcal{P}$ is a refinement of $\mathcal{R}$ and a coarsening of $\mathcal{B}$. Moreover, every $\mathcal{P}_{\mathcal{D}}$ is a coarsening of $\mathcal{B}_{\mathcal{D}}$. Since $\bigcup \mathcal{B}_{\mathcal{D}} = \bigcup_{C \in \mathcal{D}} \partial_G C$, every $A \in \mathcal{P}$ is contained in the union of the boundaries of components in some $\mathcal{D}_A \in \mathfrak{R}$.

For every $A \in \mathcal{P}$, we set $R_A := L'_{\mathcal{D}_A} + \ell + K$ and $r_A := L_{\mathcal{D}_A}$. Note that $2\ell + 3K = r' + \ell + K \leq R_A \leq L'(t, K)$ and $0 < r \leq r_A \leq \ell$. This completes the construction at step $n + 1$.

It remains to check that every $A \in \mathcal{P}$ has diameter at most $R_0(t, K)$ and that $\mathcal{A}$ and the R_A, r_A satisfy (1) to (3). For every $A \in \mathcal{A}$ we have

$$\operatorname{diam}_G(A) \leq n'D' + (n' - 1)(2r' + n'd') \leq R_0(t, K)$$

by Lemma 6.1 (iii).

To prove (1), let $A \in \mathcal{P}$ and $h := h_A$. By the choice of R_h, r_h we have $R_h = L'_{\mathcal{D}_A} + \ell + K$ and $r_h = L_{\mathcal{D}_A}$, and hence (1) follows from (1″).

To prove (3), let C be a component of $G - G^n$. Since $\mathcal{P}_{\mathcal{D}}$ is a coarsening of $\mathcal{B}_{\mathcal{D}}$ and $\mathcal{P}$ is the union over all $\mathcal{P}_{\mathcal{D}}$ with $\mathcal{D} \in \mathfrak{R}$, there are at most m_C elements of $\mathcal{P}$ that meet C. By the definition of the new partition classes V_{h_A} as $B_{G-G^n}(A, R_A)$, it follows that at most m_C partition classes of $\mathcal{H}^{n+1}$ meet C. Since $m_C \leq t - 1$ as shown earlier, this concludes the proof of (3).

To prove (2), let $g \neq h \in V(H^{n+1})$ be non-adjacent. By (2) of H^n, it suffices to consider the case where $g \in L^{n+1} = V(H^{n+1} - H^n)$. If $h \in V(H^{n-1})$, then

$$\begin{aligned} d_G(V_g, V_h) &\geq d_G(G - G^n, G^{n-1}) \geq \min\{R_{h'} : h' \in L^n\} \\ &\geq 2\ell + 3K \geq 2 \cdot \max\{r_g, r_h\} + 3K. \end{aligned} \tag{a}$$

where the second inequality holds by the definition of the $V_{h'}$, the third inequality holds because $R_{h'} \geq 2\ell + 3K$, and the last inequality holds because $r_g, r_h \leq \ell$.

Now assume $h \in L^n = V(H^n - H^{n-1})$, and let P be a V_g–V_h path in G of length $d_G(V_g, V_h)$. As $gh \notin E(H^{n+1})$ and L^{n+1} is independent, P meets either $G - G^{n+1}$ or it meets a bag $V_{h'}$ of some $h' \neq h \in L^n$ (see Figure 7). In the former case, we obtain $d_G(V_g, V_h) \geq d_G(G - G^{n+1}, G^n) \geq 2 \cdot \max\{r_g, r_h\} + 3K$ by the same argument as in (a). In the latter case, since L^n is independent, P has to meet either G^{n-1}, and we are done as before, or P meets a bag $V_{g'}$ of some $g' \neq g \in L^{n+1}$ (see Figure 7). Then $d_G(V_g, V_h) \geq \max\{d_G(V_g, V_{g'}), d_G(V_h, V_{h'})\} \geq 2 \cdot \max\{r_g, r_{g'}, r_h, r_{h'}\} + 3K$ by (2), once we have proved that (2) holds for $g, g' \in L^{n+1}$.

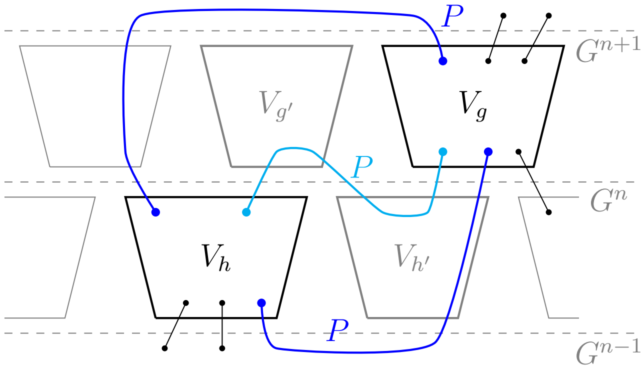

Figure 7: Depicted is the case where $g \in L^{n+1}$ and $h \in L^n$. The light blue path meets both $V_{h'}$ and $V_{g'}$. The dark blue paths meet either $G - G^{n+1}$ or G^{n-1}.

Hence, it remains to consider the case where $g \neq h \in L^{n+1}$, i.e. $g = h_A$ and $h = h_B$ for some $A, B \in \mathcal{P}$. Let us first assume that $\mathcal{D}_A \neq \mathcal{D}_B$. If at least one of $\mathcal{D}_A, \mathcal{D}_B$ is of the form $\{C\}$ for some $C \in \mathcal{C}(G - G^n)$, then V_{h_A}, V_{h_B} can be separated in G by the partition class V_x of the (unique) node $h \in L^n$ with $N_G(C) \subseteq V_h$. Otherwise, V_{h_A}, V_{h_B} are contained in distinct components of $G - G^{n-1}$. Hence, any V_{h_A}–V_{h_B} path meets G^{n-1}, and so $d_G(V_{h_A}, V_{h_B}) \geq d_G(G - G^n, G^{n-1}) \geq 2 \cdot \max\{r_{h_A}, r_{h_B}\} + 3K$ as in (a).

Thus, we may assume $\mathcal{D}_A = \mathcal{D}_B$. Let P be a V_{h_A}–V_{h_B} path in G of length $d_G(V_{h_A}, V_{h_B})$. If P has a subpath that lies in $G - G^n$ and starts in V_{h_A} and ends in $V_{h_{B'}}$ for some $B' \neq A \in \mathcal{P}$ with $\mathcal{D}_{B'} = \mathcal{D}_A$, then

$$d_G(V_{h_A}, V_{h_B}) \geq d_{G-G^n}(V_{h_A}, V_{h_{B'}}) \geq d_{G-G^n}(A, B') - R_{h_A} - R_{h_{B'}}$$

and hence, by ($\triangle$) and the definition of R_{h_A}, R_{h_B},

$$\begin{aligned} d_G(V_{h_A}, V_{h_B}) &\geq (2L'_{\mathcal{D}_A} + 4\ell + 5K) - 2 \cdot (L'_{\mathcal{D}_A} + \ell + K) \\ &\geq 2\ell + 3K \geq 2 \cdot \max\{r_{h_A}, r_{h_B}\} + 3K. \end{aligned}$$

Otherwise, since L^{n+1} is independent, P has a subpath that starts in A and ends in B' for some $B' \in \mathcal{P}$ with $\mathcal{D}_{B'} = \mathcal{D}_A$. Then by ($\square$)

$$d_G(V_{h_A}, V_{h_B}) \geq d_G(A, B') \geq 2L_{\mathcal{D}_A} + 3K = 2 \cdot \max\{r_{h_A}, r_{h_B}\} + 3K.$$

This establishes (2) and hence concludes the proof. □

8 The approximation algorithm

Note that our proof of Lemma 3.4 is constructive (and so are any lemmas it relies on), and therefore we will be able to turn it into an algorithm that approximates, to a constant factor, the optimal distortion $\alpha_t(G)$ of any embedding of a finite graph G into a $K_{2,t}$-minor-free graph in polynomial time, thereby proving Corollary 1.3:

Proof of Corollary 1.3. Let $n := |V(G)|$. For each $K = 1, 2, \ldots n$, our algorithm attempts to carry out the construction of H and the H-partition of G as in the proof of Lemma 3.4, without knowing in advance whether G has a K-fat $K_{2,t}$ minor. Note that the only occasions in that proof where we used the assumption that G has no such minor were when invoking Lemmas 5.2 and 5.3. Thus, either the attempt will output such an H, or one of these calls to the aforementioned Lemmas will return a K-fat $K_{2,t}$ minor model in G, in which case we say that the attempt failed. In the former case, where our algorithm constructs a graph H and an H-partition of G, it then checks whether H is $K_{2,t}$-minor-free (which can be done in polynomial time [17]). If H is $K_{2,t}$-minor-free, then we say that the attempt was *successful*. If not, then the attempt failed, and invoking Corollary 4.1 again returns a K-fat model of $K_{2,t}$ in G.

Our algorithm returns the smallest value $K_{\min}$ of $K \leq n$ for which this procedure succeeds as an approximate value for $\alpha_t(G)$. Note that $K_{\min}$ exists since G cannot have a n-fat $K_{2,0}$ minor. Along with $K_{\min}$, the algorithm can return a witness: we start with the graph H and the embedding of G into H, defined by mapping each $v \in V(G)$ into its partition class $V_h \ni v$, and then modify H and the embedding using the star trick mentioned before the statement of Corollary 1.3 to eliminate the additive error.

We claim that $K_{\min}$ is within a constant factor of $\alpha_t(G)$. Indeed, our Theorem 1.2 (and the remark thereafter) guarantees that the multiplicative distortion of G into H, which is $K_{2,t}$-minor-free by definition, is at most $C \cdot K_{\min}$ for a universal constant C. If $K_{\min} > 1$ then our procedure failed for $K = K_{\min} - 1$, and therefore as mentioned above it will identify a $(K-1)$-fat $K_{2,t}$ minor model M in G. It is not hard to see that such a model implies that $\alpha_t(G)$ is at least $c \cdot (K_{\min} - 1)$ for a small universal constant c [9, Proposition 3] (the precise value of which depends on the convention chosen in the definition of multiplicative distortion). Our algorithm outputs M as a witness for this lower bound on $\alpha_t(G)$. If on the other hand $K_{\min} = 1$, then as above we deduce that $\alpha_t(G) \leq C$, and so we do not need a lower bound or a witness, as we can use the trivial bound $\alpha_t(G) \geq 1$. □

Note that our algorithm has the pleasant property that if $\alpha_t(G) = 1$, equivalently, if G is $K_{2,t}$-minor-free, then the output of our algorithm is 1.

Both the running time of our algorithm, and the approximation constant we obtained, increase with t. We do not know to what extent this is necessary.

Our Corollary 1.3, along with analogous results of [1], and remarks of [9], motivates the following problem related to the coarse Menger conjecture of [5, 13]. Given a finite graph G, and $S, T \subset V(G)$, and $n \in \mathbb{N}$, let $MM_n(G, S, T)$ denote the maximum $K \in \mathbb{N}$ such that there is an n-tuple of S–T paths in G pairwise at distance at least K.

Problem 8.1. *Is it true that for every $n \geq 2$, there are universal constants $C, c > 1$, such that:*

(i) *there is an efficient algorithm that, given G, S, T as above, approximates $MM_n(G, S, T)$ up to a multiplicative factor of C; and*

(ii) *approximating $MM_n(G, S, T)$ up to a multiplicative factor of c is NP-hard.*

We remark that the results of [5, 13] that the coarse Menger conjecture is true for $n = 2$ imply that (i) holds for $n = 2$: the algorithm can output the smallest radius of a ball in G separating S from T. This trivially lower-bounds $MM_2(G, S, T)$, and the aforementioned result states that it is also an upper bound up to a universal constant C.

If we require the exact rather than an approximate value for $MM_n(G, S, T)$, the problem is NP-hard as proved by Baligács and MacManus[7].

We do not know whether the analogue of (ii) holds for $\alpha_t(G)$.

References

[1] R. Agarwala, V. Bafna, M. Farach, B. Narayanan, M. Paterson, and M. Thorup. On the approximability of numerical taxonomy (fitting distances by tree metrics). *SIAM J. Comput.*, 28:1073–1085, 1999.

[2] S. Albrechtsen and J. Davies. Counterexample to the conjectured coarse grid theorem. arXiv:2508.15342.

[3] S. Albrechtsen, R. Diestel, A.-K. Elm, E. Fluck, R. W. Jacobs, P. Knappe, and P. Wollan. A structural duality for path-decompositions into parts of small radius. arXiv:2307.08497.

[4] S. Albrechtsen, M. Distel, and A. Georgakopoulos. Small counterexamples to the fat minor conjecture. In preparation.

[5] S. Albrechtsen, T. Huynh, R. W. Jacobs, P. Knappe, and P. Wollan. A Menger-Type Theorem for Two Induced Paths. *SIAM J. Discrete Math.*, 38(2):1438–1450, 2024.

[6] S. Albrechtsen, R. W. Jacobs, P. Knappe, and P. Wollan. A characterisation of graphs quasi-isometric to K_4-minor-free graphs. arXiv:2408.15335.

[7] J. Baligács and J. MacManus. The metric Menger problem. arXiv:2403.05630.

[8] M. Bonamy, N. Bousquet, L. Esperet, C. Groenland, C.-H. Liu, F. Pirot, and A. Scott. Asymptotic Dimension of Minor-Closed Families and Assouad-Nagata Dimension of Surfaces. *J. Eur. Math. Soc.*, 26(10):3739–3791, 2023.

[9] V. Chepoi, F. F. Dragan, I. Newman, Y. Rabinovich, and Y. Vaxès. Constant Approximation Algorithms for Embedding Graph Metrics into Trees and Outerplanar Graphs. *Discrete & Computational Geometry*, 47(1):187–214, 2012.

[10] J. Davies, R. Hickingbotham, F. Illingworth, and R. McCarty. Fat minors cannot be thinned (by quasi-isometries). arXiv:2405.09383.

[11] R. Diestel. *Graph Theory* (6th edition). Springer-Verlag, 2024.

[12] K. Fujiwara and P. Papasoglu. A coarse-geometry characterization of cacti. arXiv:2305.08512.

[13] A. Georgakopoulos and P. Papasoglu. Graph minors and metric spaces. *Combinatorica*, 45:33, 2025.

[14] T. Nguyen, A. Scott, and P. Seymour. Asymptotic structure. II. Path-width and additive quasi-isometry. Preprint 2024.

[15] T. Nguyen, A. Scott, and P. Seymour. Asymptotic structure. IV. A counterexample to the weak coarse Menger conjecture. arXiv:2508.14332.

[16] T. Nguyen, A. Scott, and P. Seymour. Coarse tree-width. arXiv:2501.09839.

[17] N. Roberston and P. Seymour. Graph minors. XIII. The disjoint paths problem. *J. Combin. Theory (Series B)*, 63(1):65–110, 1995.