



HAL
open science

Packing and Covering Balls in Graphs Excluding a Minor

Nicolas Bousquet, Wouter Cames van Batenburg, Louis Esperet, Gwenaël Joret, William Lochet, Carole Muller, François Pirot

► **To cite this version:**

Nicolas Bousquet, Wouter Cames van Batenburg, Louis Esperet, Gwenaël Joret, William Lochet, et al.. Packing and Covering Balls in Graphs Excluding a Minor. *Combinatorica*, 2021, 41 (3), pp.299-318. 10.1007/s00493-020-4423-3 . hal-03281264

HAL Id: hal-03281264

<https://hal.science/hal-03281264v1>

Submitted on 12 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

PACKING AND COVERING BALLS IN GRAPHS EXCLUDING A MINOR

NICOLAS BOUSQUET, WOUTER CAMES VAN BATENBURG, LOUIS ESPERET, GWENAËL JORET,
WILLIAM LOCHET, CAROLE MULLER, AND FRANÇOIS PIROT

ABSTRACT. We prove that for every integer $t \geq 1$ there exists a constant c_t such that for every K_t -minor-free graph G , and every set S of balls in G , the minimum size of a set of vertices of G intersecting all the balls of S is at most c_t times the maximum number of vertex-disjoint balls in S . This was conjectured by Chepoi, Estellon, and Vaxès in 2007 in the special case of planar graphs and of balls having the same radius.

1. INTRODUCTION

A hypergraph \mathcal{H} is a pair (V, \mathcal{E}) where V is the vertex set and $\mathcal{E} \subseteq 2^V$ is the edge set of \mathcal{H} . A *matching* in a hypergraph \mathcal{H} is a set of pairwise vertex-disjoint edges, and a *transversal* is a set of vertices that intersects every edge. The *matching number* of a hypergraph \mathcal{H} , denoted by $\nu(\mathcal{H})$, is the maximum number of edges in a matching. The *transversal number* of \mathcal{H} , denoted by $\tau(\mathcal{H})$, is the minimum size of a transversal of \mathcal{H} . We can also consider the linear relaxation of these two parameters: we define the *fractional matching number* $\nu^*(\mathcal{H})$ and the *fractional transversal number* $\tau^*(\mathcal{H})$ as follows.

$$\nu^*(\mathcal{H}) = \max \sum_{e \in \mathcal{E}(\mathcal{H})} w_e$$

given that $\begin{cases} \sum_{e \ni v} w_e \leq 1 & \text{for every vertex } v \text{ of } \mathcal{H} \\ w_e \geq 0 & \text{for every edge } e \text{ of } \mathcal{H}, \end{cases}$

and the dual of this linear program is

$$\tau^*(\mathcal{H}) = \min \sum_{v \in V(\mathcal{H})} w_v$$

given that $\begin{cases} \sum_{v \in e} w_v \geq 1 & \text{for every edge } e \text{ of } \mathcal{H} \\ w_v \geq 0 & \text{for every vertex } v \text{ of } \mathcal{H}. \end{cases}$

By the strong duality theorem, $\nu(\mathcal{H}) \leq \nu^*(\mathcal{H}) = \tau^*(\mathcal{H}) \leq \tau(\mathcal{H})$ for every hypergraph \mathcal{H} . Given a class \mathcal{C} of hypergraphs, a classical problem in combinatorial optimization is to decide whether there exists a function f such that $\tau(\mathcal{H}) \leq f(\nu(\mathcal{H}))$ for every $\mathcal{H} \in \mathcal{C}$. If this is the case the class \mathcal{C} is sometimes said to have the *Erdős-Pósa property*. Classical examples include the family of all cycles of a graph [12] (i.e. given a graph $G = (V, E)$ we consider the hypergraph with vertex set V whose edges are all the cycles of G), and the family of all

N. Bousquet is supported by ANR Project DISTANCIA (ANR-17-CE40-0015). W. Cames van Batenburg, G. Joret, and F. Pirot are supported by an ARC grant from the Wallonia-Brussels Federation of Belgium. L. Esperet is supported by ANR Projects GATO (ANR-16-CE40-0009-01) and GrR (ANR-18-CE40-0032). C. Muller is supported by the Luxembourg National Research Fund (FNR) (Grant Agreement Nr 11628910).

directed cycles of a directed graph [23]. A desirable property is the existence of a constant c such that $\tau(\mathcal{H}) \leq c \cdot \tau^*(\mathcal{H})$ or $\nu^*(\mathcal{H}) \leq c \cdot \nu(\mathcal{H})$, or (even better) $\tau(\mathcal{H}) \leq c \cdot \nu(\mathcal{H})$ for every $\mathcal{H} \in \mathcal{C}$. These properties are often useful in the design of approximation algorithms using a primal-dual approach (see for instance [16, 14]).

Given a graph $G = (V, E)$, an integer $r \geq 0$, and a vertex $v \in V$, we denote by $B_r(v)$ the *ball of radius r* in G centered in v , that is

$$B_r(v) := \{u \in V(G) \mid d_G(u, v) \leq r\},$$

where $d_G(u, v)$ denotes the distance between u and v in G (we will omit the subscript G when the graph is clear from the context). We say that a hypergraph \mathcal{H} is a *ball hypergraph* of G if \mathcal{H} has vertex set $V = V(G)$ and each edge of \mathcal{H} is a ball $B_r(v)$ in G for some integer r and some vertex $v \in V$. If all the balls forming the edges of \mathcal{H} have the same radius r , we say that \mathcal{H} is an *r -ball hypergraph* of G .

Chepoi, Estellon, and Vaxès [7] proved the existence of a universal constant ρ such that for every $r \geq 0$ and every planar graph G of diameter at most $2r$, the vertices of G can be covered with at most ρ balls of radius r . This result was extended to graphs embeddable on a fixed surface with a bounded number of apices in [1]. Note that G has diameter at most $2r$ if and only if there are no two disjoint balls of radius r in G . Also, a set of balls of radius r in G covers all of $V(G)$ if and only if their centers intersect all balls of radius r in G . Thus, these results state equivalently the existence of a universal constant ρ such that for every $r \geq 0$ and every planar (or more generally bounded genus) graph G , if the r -ball hypergraph \mathcal{H} consisting of all balls of radius r satisfies $\nu(\mathcal{H}) = 1$, then $\tau(\mathcal{H}) \leq \rho$. With this interpretation in mind, Chepoi, Estellon, and Vaxès [3] conjectured the following generalization in 2007 (see also [13]).

Conjecture 1 (Chepoi, Estellon, and Vaxès [3]). *There exists a constant c such that for every integer $r \geq 0$, every planar graph G , and every r -ball hypergraph \mathcal{H} of G , we have $\tau(\mathcal{H}) \leq c \cdot \nu(\mathcal{H})$.*

If one considers all metric spaces obtained as standard graph-metrics of planar graphs, then Conjecture 1 states that these metric spaces satisfy the so-called *bounded covering-packing property* [6]. Recently, Chepoi, Estellon, and Naves [6] showed that other metric spaces do have this property, including the important case of Busemann surfaces. (Quoting [6], the latter are roughly the geodesic metric spaces homeomorphic to \mathbb{R}^2 in which the distance function is convex; they generalize Euclidean spaces, hyperbolic spaces, Riemannian manifolds of global nonpositive sectional curvatures, and CAT(0) spaces.)

Going back to Conjecture 1, let us emphasize that a key aspect of this conjecture is that the constant c is independent of the radius r . If c is allowed to depend on r , then the conjecture is known to be true. In fact, it holds more generally for all graph classes with bounded expansion, as shown by Dvořák [10].

Some evidence for Conjecture 1 was given by Bousquet and Thomassé [2], who proved that it holds with a polynomial bound instead of a linear one. More generally, they proved that for every integer $t \geq 1$, there exists a constant c_t such that for every integer $r \geq 0$, every K_t -minor free graph G , and every r -ball hypergraph \mathcal{H} of G , we have $\tau(\mathcal{H}) \leq c_t \cdot \nu(\mathcal{H})^{2t+1}$.

The main result of this paper is that Conjecture 1 is true, and furthermore it is not necessary to assume that all the balls have the same radius.

Theorem 2 (Main result). *For every integer $t \geq 1$, there is a constant c_t such that $\tau(\mathcal{H}) \leq c_t \cdot \nu(\mathcal{H})$ for every K_t -minor-free graph G and every ball hypergraph \mathcal{H} of G .*

A set S of vertices of a graph G is r -dominating if each vertex of G is at distance at most r from S , and r -independent if any two vertices of S are at distance at least $2r + 1$ apart in G . Note that if we take \mathcal{H} to be the r -ball hypergraph consisting of all balls of radius r in G , Theorem 2 has the following interesting graph-theoretic interpretation: if G is K_t -minor-free, then the minimum size of an r -dominating set is at most c_t times the maximum size of an r -independent set in G .

Our proof of Theorem 2 follows a bootstrapping approach. It relies on the existence of some function f_t such that $\tau(\mathcal{H}) \leq f_t(\nu(\mathcal{H}))$, i.e. on the Erdős-Pósa property of the ball hypergraphs of K_t -minor-free graphs, which is used in the proof when $\nu(\mathcal{H})$ is not ‘too big’. However, showing this property was an open problem. This was known for r -ball hypergraphs, by the result of Bousquet and Thomassé [2], but their proof method does not extend to the case of balls of arbitrary radii. For this reason, as a first step towards proving Theorem 2, we prove Theorem 3 below establishing said Erdős-Pósa property. We also note that, while the bounding function in Theorem 3 is not optimal, it is a near linear bound of the form $\tau(\mathcal{H}) \leq c_t \cdot \nu(\mathcal{H}) \log \nu(\mathcal{H})$ where c_t is a small explicit constant polynomial in t . This is in contrast with the constant c_t in our proof of Theorem 2 which is large, exponential in t . Thus, the bound in Theorem 3 is better for small values of $\nu(\mathcal{H})$. (We note that logarithms in this paper are natural, and the base of the natural logarithm is denoted by e .)

Theorem 3 (Near linear bound). *Let G be a graph with no K_t -minor and such that every minor of G has average degree at most d . Then for every ball hypergraph \mathcal{H} of G ,*

$$\tau(\mathcal{H}) \leq 2e(t-1)d \cdot \nu(\mathcal{H}) \cdot \log(11ed \cdot \nu(\mathcal{H})).$$

In particular, $\tau(\mathcal{H}) \leq ct^2\sqrt{\log t} \cdot \nu(\mathcal{H}) \cdot \log(t \cdot \nu(\mathcal{H}))$ for some absolute constant $c > 0$, and if G is planar then $\tau(\mathcal{H}) \leq 48e \cdot \nu(\mathcal{H}) \cdot \log(66e \cdot \nu(\mathcal{H}))$.

The proof of Theorem 3 uses known results on the VC-dimension of ball hypergraphs of G when G excludes a minor, together with classical bounds relating $\tau(\mathcal{H})$ and $\tau^*(\mathcal{H})$ when \mathcal{H} has bounded VC-dimension, as well as the following theorem.

Theorem 4 (Fractional version). *Let G be a graph and let d be the maximum average degree of a minor of G . Then for every ball hypergraph \mathcal{H} of G , we have $\nu^*(\mathcal{H}) \leq ed \cdot \nu(\mathcal{H})$. In particular, if G is planar then $\nu^*(\mathcal{H}) \leq 6e \cdot \nu(\mathcal{H})$ and if G has no K_t -minor then $\nu^*(\mathcal{H}) \leq c \cdot t\sqrt{\log t} \cdot \nu(\mathcal{H})$, for some absolute constant $c > 0$.*

We note that results on the VC-dimension of ball hypergraphs in graphs excluding a minor have also been used recently to obtain improved algorithms for the computation of the diameter in sparse graphs [9, 20].

The proofs of Theorems 2, 3, and 4 are constructive, and can be transformed into efficient algorithms producing transversals (in the case of Theorems 2 and 3) or matchings (in the case of Theorem 4) of the desired size.

The paper is organized as follows. Sections 2 and 3 are devoted to technical lemmas that will be used in our proofs. Theorems 3 and 4 are proved in Section 4. Theorem 2 is proved in Section 5. Finally, we conclude the paper in Section 6 with a construction suggesting that Theorem 2 does not extend way beyond proper minor-closed classes.

2. HYPERGRAPHS, BALLS, AND MINORS

We will need two technical lemmas, whose proofs are very similar to the proof of [2, Theorem 4] and [7, Proposition 1]. We start with Lemma 1, which will be used in the proofs of Theorem 4 and Theorem 2. We first need the following definitions.

We say that two balls B_1 and B_2 are *comparable* if $B_1 \subset B_2$ or $B_2 \subset B_1$ (where \subset denotes the strict inclusion), and otherwise they are *incomparable*. Note in particular that if $B_1 = B_2$ then B_1 and B_2 are incomparable. Consider two intersecting and incomparable balls $B_1 = B_{r_1}(v_1)$ and $B_2 = B_{r_2}(v_2)$ in a graph G , and let $d := d_G(v_1, v_2)$. A *median vertex* of B_1 and B_2 is any vertex u lying on a shortest path between v_1 and v_2 , at distance $\lfloor \frac{r_1 - r_2 + d}{2} \rfloor$ from v_1 and at distance $\lceil \frac{r_2 - r_1 + d}{2} \rceil$ from v_2 , or symmetrically at distance $\lceil \frac{r_1 - r_2 + d}{2} \rceil$ from v_1 and at distance $\lfloor \frac{r_2 - r_1 + d}{2} \rfloor$ from v_2 . Since B_1 and B_2 intersect, we have $r_1 + r_2 \geq d$ and since B_1 and B_2 are incomparable, we have $r_2 \leq r_1 + d$ and $r_1 \leq r_2 + d$, and in particular $\lfloor \frac{r_1 - r_2 + d}{2} \rfloor \geq 0$ and $\lceil \frac{r_2 - r_1 + d}{2} \rceil \geq 0$ (so the distances above are well defined). Moreover, $\lfloor \frac{r_1 - r_2 + d}{2} \rfloor = \lfloor \frac{2r_1 - r_1 - r_2 + d}{2} \rfloor \leq r_1$ and $\lceil \frac{r_2 - r_1 + d}{2} \rceil = \lceil \frac{2r_2 - r_1 - r_2 + d}{2} \rceil \leq r_2$, so any median vertex of B_1 and B_2 lies in $B_1 \cap B_2$. Finally, note that by the definition of a median vertex u of B_1 and B_2 ,

- for every $\{i, j\} = \{1, 2\}$ we have $r_j - d(v_j, u) \leq r_i - d(v_i, u) + 1$, and
- if $v_1 = v_2$ (which implies $B_1 = B_2$ since the balls are incomparable), then $u = v_1 = v_2$.

Lemma 1. *Let G be a graph, let $S = \{B_i = B_{r_i}(s_i)\}_{i \in [n]}$ be a set of n pairwise incomparable balls in G , with pairwise distinct centers, and let $E_S \subseteq \binom{S}{2}$ be a subset of pairs of intersecting balls $\{B_i, B_j\} \subseteq S$, each of which is associated with a median vertex $x_{\{i,j\}}$ of B_i and B_j , and such that the only balls of S containing $x_{\{i,j\}}$ are B_i and B_j . Then the graph $H = (S, E_S)$ is a minor of G .*

Proof. Let us fix a total ordering \prec on the vertices of G . In the proof, all distances are in the graph G , so we write $d(u, v)$ instead of $d_G(u, v)$ for the sake of readability. For every pair of balls $\{B_i, B_j\} \in E_S$, we write x_{ij} or x_{ji} instead of $x_{\{i,j\}}$, for the sake of readability (x_{ij} , x_{ji} , and $x_{\{i,j\}}$ all correspond to the same median vertex of B_i and B_j). We also let $P(s_i, x_{ij})$ be a shortest path from s_i to x_{ij} , and we assume that the sequence of vertices from s_i to x_{ij} on the path is minimum with respect to the lexicographic order induced by \prec (among all shortest paths from s_i to x_{ij}). By the assumptions, we know that $P_{ij} := P(s_i, x_{ij}) \cup P(s_j, x_{ij})$ is a shortest path from s_i to s_j .

For every $i \in [n]$, we define

$$\mathcal{T}_i := \bigcup_{j: \{B_i, B_j\} \in E_S} P(s_i, x_{ij}).$$

Claim 1. For every $i \in [n]$, \mathcal{T}_i is a tree.

Assume for the sake of contradiction that there is a cycle C in \mathcal{T}_i . Observe that, by construction, if uv is an edge of \mathcal{T}_i then $|d(s_i, u) - d(s_i, v)| = 1$. Let y be a vertex of C maximizing $d(s_i, y)$, and let z_1, z_2 denote its two neighbors in C . Then $d(s_i, z_1) = d(s_i, z_2) = d(s_i, y) - 1$, and there exist j_1, j_2 such that z_1y is an edge of $P(s_i, x_{ij_1})$ and z_2y is an edge of $P(s_i, x_{ij_2})$. Let P_1 and P_2 be the subpaths from s_i to y of $P(s_i, x_{ij_1})$ and $P(s_i, x_{ij_2})$, respectively. Then P_1 and P_2 are two different paths from s_i to y , and one of them is not minimum either in terms of length, or with respect to the lexicographic order induced by \prec . This contradicts the definition of $P(s_i, x_{ij_1})$ and $P(s_i, x_{ij_2})$.

Claim 2. For every two pairs of balls $\{B_i, B_k\}, \{B_j, B_\ell\} \in E_S$ with $i \neq j$, if $P(s_i, x_{ik})$ and $P(s_j, x_{j\ell})$ intersect in some vertex y such that $d(y, x_{ik}) \leq d(y, x_{j\ell})$, then $j = k$ and $y = x_{ij}$.

Note that $d(s_j, x_{ik}) \leq d(s_j, y) + d(y, x_{ik}) \leq d(s_j, y) + d(y, x_{j\ell}) = d(s_j, x_{j\ell})$. Since $x_{j\ell}$ is a median vertex of B_j and B_ℓ , we have $d(s_j, x_{j\ell}) \leq r_j$, which implies that $d(s_j, x_{ik}) \leq r_j$ and

thus $x_{ik} \in B_j$. By definition, x_{ik} is only contained in the balls B_i and B_k of S and thus $j = k$. If we also have $i = \ell$, then necessarily $y = x_{ij}$.

From now on, we assume that $i \neq \ell$. Since $P_{ij} = P(s_i, x_{ij}) \cup P(s_j, x_{ij})$ is a shortest path containing the vertex y , the s_j - y section of that path (which contains x_{ij}) has the same length as the s_j - y section of $P(s_j, x_{j\ell})$. Replacing the latter section by the former, we obtain a shortest path from s_j to $x_{j\ell}$ containing x_{ij} , which we denote $Q(s_j, x_{j\ell})$. As a consequence,

$$d(x_{j\ell}, x_{ij}) = d(x_{j\ell}, s_j) - d(s_j, x_{ij}) \leq r_j - d(s_j, x_{ij}) \leq r_i - d(s_i, x_{ij}) + 1,$$

where the last inequality follows from the definition of x_{ij} . We now use the fact that y appears on the path $P(s_i, x_{ij})$ and on the x_{ij} - $x_{j\ell}$ section of $Q(s_j, x_{j\ell})$, and obtain

$$d(s_i, x_{j\ell}) \leq d(s_i, y) + d(y, x_{j\ell}) = d(s_i, x_{ij}) + d(x_{ij}, x_{j\ell}) - 2d(y, x_{ij}) \leq r_i + 1 - 2d(y, x_{ij}).$$

Since $x_{j\ell} \notin B_i$ by definition (and so $d(s_i, x_{j\ell}) > r_i$), this implies that $y = x_{ij}$, as desired.

This claim immediately implies that for every $i, j \in [n]$ with $i \neq j$, we have $V(\mathcal{T}_i) \cap V(\mathcal{T}_j) = \{x_{ij}\}$ if $\{B_i, B_j\} \in E_S$, and $V(\mathcal{T}_i) \cap V(\mathcal{T}_j) = \emptyset$ otherwise. Another consequence is that for every $\{B_i, B_j\} \in E_S$, the vertex x_{ij} is a leaf in at least one of the two trees \mathcal{T}_i and \mathcal{T}_j (since otherwise there exist $k \neq j$ and $\ell \neq i$ such that $x_{ij} \in P(s_i, x_{ik})$ and $x_{ij} \in P(s_j, x_{j\ell})$, which readily contradicts Claim 2 above).

In the subgraph $\bigcup_{i \in [n]} \mathcal{T}_i$ of G , for each $i \in [n]$ we contract each edge of \mathcal{T}_i except the ones incident to a leaf of \mathcal{T}_i . It follows from the paragraph above that the resulting graph is precisely a graph obtained from $H = (S, E_S)$ by subdividing each edge at most once, and thus H is a minor of G . \square

The next result has a very similar proof¹, but the setting is slightly different. It will be used in the proof of Theorem 2.

Lemma 2. *Let G be a graph and $S = \{B_i = B_{r_i}(s_i)\}_{i \in [n]}$ be a set of n pairwise vertex-disjoint balls in G , and let $E_S \subseteq \binom{S}{2}$ be a subset of pairs of balls $\{B_i, B_j\} \subseteq S$, each of which is associated with a ball $B_{\{i,j\}} \notin S$ of G which intersects only B_i and B_j in S . Then the graph $H = (S, E_S)$ is a minor of G .*

Proof. Let us fix a total ordering \prec on the vertices of G . As before, all distances are in the graph G , and we write $d(u, v)$ instead of $d_G(u, v)$. For every $\{B_i, B_j\} \in E_S$ we write B_{ij} or B_{ji} interchangeably for $B_{\{i,j\}}$, and we denote by x_{ij} the center of the ball B_{ij} , and by r_{ij} its radius ($x_{ij} = x_{ji}$ and $r_{ij} = r_{ji}$). We can assume that the centers x_{ij} are chosen so that the radii r_{ij} are minimal (among all balls of G not in S that intersect only B_i and B_j in S).

We let $P(s_i, x_{ij})$ be the shortest path from s_i to x_{ij} which minimizes the sequence of vertices from s_i to x_{ij} with respect to the lexicographic ordering induced by \prec (among all shortest paths from s_i to x_{ij}). Observe that $P(s_i, x_{ij})$ and $P(s_j, x_{ij})$ only intersect in x_{ij} (if not, we could replace x_{ij} by a vertex that is closer to s_i and s_j and reduce the radius r_{ij} accordingly – the new ball B_{ij} would still intersect B_i and B_j , and no other ball of S , and this would contradict the minimality of r_{ij}). We may also assume that $r_i + r_{ij} - 1 \leq d(s_i, x_{ij}) \leq r_i + r_{ij}$ (otherwise we could replace x_{ij} by its neighbor on $P(s_j, x_{ij})$ and decrease r_{ij} by 1).

For every $i \in [n]$, we define

$$\mathcal{T}_i := \bigcup_{j: \{B_i, B_j\} \in E_S} P(s_i, x_{ij}).$$

¹Despite our best effort, we have not been able to prove the two results at once in a satisfactory way, i.e. with a proof that would be both readable and shorter than the concatenation of the two existing proofs.

Claim 1. For every $i \in [n]$, \mathcal{T}_i is a tree.

The proof is exactly the same as that of Claim 1 in the proof of Lemma 1 (we do not repeat it here).

On the path $P(s_i, x_{ij})$, we let $z_{i,ij}$ be the vertex at distance r_i from s_i (and since $x_{ij} = x_{ji}$ we use $z_{i,ij}$ and $z_{i,ji}$ interchangeably). As we assumed above that $r_i + r_{ij} - 1 \leq d(s_i, x_{ij}) \leq r_i + r_{ij}$, we also have $r_{ij} - 1 \leq d(x_{ij}, z_{i,ij}) \leq r_{ij}$. In particular, $d(x_{ij}, z_{j,ij}) - 1 \leq d(x_{ij}, z_{i,ij}) \leq d(x_{ij}, z_{j,ij}) + 1$.

Claim 2. For every two pairs of balls $\{B_i, B_k\}, \{B_j, B_\ell\} \in E_S$, with $i \neq j$, if $P(s_i, x_{ik})$ and $P(s_j, x_{j\ell})$ intersect in some vertex y such that $d(y, z_{i,ik}) \leq d(y, z_{j,j\ell})$, then $i = \ell$ and $y = x_{ij}$.

We first argue that y appears after $z_{j,j\ell}$ when traversing $P(s_j, x_{j\ell})$ from s_j to $x_{j\ell}$. Indeed, otherwise we would have

$$d(s_j, z_{i,ik}) \leq d(s_j, y) + d(y, z_{i,ik}) \leq d(s_j, y) + d(y, z_{j,j\ell}) = d(s_j, z_{j,j\ell}) = r_j,$$

which means that B_i and B_j intersect, contradicting the assumptions that $i \neq j$ and all balls in S are vertex-disjoint. So y lies on the $z_{j,j\ell}$ - $x_{j\ell}$ section of $P(s_j, x_{j\ell})$, and we infer that

$$d(x_{j\ell}, z_{i,ik}) \leq d(x_{j\ell}, y) + d(y, z_{i,ik}) \leq d(x_{j\ell}, y) + d(y, z_{j,j\ell}) = d(x_{j\ell}, z_{j,j\ell}) \leq r_{j\ell}.$$

It follows that the ball $B_{j\ell}$ intersects the ball B_i . By the assumption, this means that $i = \ell$, and thus $s_\ell = s_i$ and $z_{j,j\ell} = z_{j,ij}$. We now argue that y lies in the $z_{i,ik}$ - x_{ik} section of $P(s_i, x_{ik})$. Suppose for a contradiction that y appears strictly before $z_{i,ik}$ when traversing $P(s_i, x_{ik})$ from s_i to x_{ik} . By definition of $z_{i,ik}$, it then follows that $d(s_i, y) \leq r_i - 1$. On the other hand

$$d(s_j, y) = d(s_j, x_{ij}) - d(y, x_{ij}) = d(s_j, z_{j,ij}) + d(z_{j,ij}, x_{ij}) - d(y, x_{ij}).$$

Note that $r_i + r_{ij} - 1 \leq d(s_i, x_{ij}) \leq d(s_i, y) + d(y, x_{ij}) \leq r_i - 1 + d(y, x_{ij})$, and thus $d(y, x_{ij}) \geq r_{ij} \geq d(z_{i,ij}, x_{ij})$. We obtain that $d(z_{j,ij}, x_{ij}) \leq d(z_{i,ij}, x_{ij}) + 1 \leq d(y, x_{ij}) + 1$, and it follows that $d(s_j, y) \leq d(s_j, z_{j,ij}) + 1 = r_j + 1$. Hence $d(s_i, s_j) \leq d(s_i, y) + d(y, s_j) \leq r_i + r_j$, so B_i and B_j intersect, a contradiction. We conclude that y lies in the $z_{i,ik}$ - x_{ik} section of $P(s_i, x_{ik})$, and thus $d(x_{ik}, y) + d(y, z_{i,ik}) = d(x_{ik}, z_{i,ik})$.

Recall that by the initial assumption of the claim, combined with $i = \ell$, we have $d(y, z_{i,ik}) \leq d(y, z_{j,ij})$. Assume first that $d(y, z_{i,ik}) = d(y, z_{j,ij})$. Then

$$d(x_{ik}, z_{j,ij}) \leq d(x_{ik}, y) + d(y, z_{j,ij}) = d(x_{ik}, y) + d(y, z_{i,ik}) \leq r_{ik},$$

which implies that B_j intersects B_{ik} . Thus $j = k$, $P(s_i, x_{ik}) = P(s_i, x_{ij})$, and $P(s_j, x_{j\ell}) = P(s_j, x_{ij})$. Since these two paths have only x_{ij} in common, in this case we conclude that $y = x_{ij}$. We can now assume that $d(y, z_{i,ik}) \leq d(y, z_{j,ij}) - 1$. Recall that by definition of x_{ij} , we have $d(x_{ij}, z_{i,ij}) \geq d(x_{ij}, z_{j,ij}) - 1$, which implies that

$$d(y, z_{i,ik}) + d(y, x_{ij}) \leq d(y, z_{j,ij}) - 1 + d(y, x_{ij}) = d(z_{j,ij}, x_{ij}) - 1 \leq d(z_{i,ij}, x_{ij}).$$

Since $z_{i,ik}$ and $z_{i,ij}$ are both at distance r_i from s_i and $P(s_i, x_{ij})$ is a shortest path from s_i to x_{ij} , it follows that the concatenation of the s_i - y section of $P(s_i, x_{ik})$ and the y - x_{ij} section of $P(s_j, x_{ij})$ is a shortest path from s_i to x_{ij} (containing y). As y is also on a shortest path from s_j to x_{ij} , if we had $d(y, x_{ij}) > 0$, then we could replace x_{ij} by y and reduce r_{ij} to $r_{ij} - d(y, x_{ij})$ (B_{ij} would still intersect B_i and B_j and only these balls of S), which would contradict the minimality of r_{ij} . It follows that $y = x_{ij}$, as desired.

As in the proof of Lemma 1, the claim implies that for $i \neq j \in [n]$, $\mathcal{T}_i \cap \mathcal{T}_j = \{x_{ij}\}$ if $\{B_i, B_j\} \in E_S$, and otherwise the trees \mathcal{T}_i and \mathcal{T}_j are vertex-disjoint. Another direct consequence is that for every $\{B_i, B_j\} \in E_S$, the vertex x_{ij} is a leaf in at least one of the two trees \mathcal{T}_i and \mathcal{T}_j . As before, we can contract the edges of each tree \mathcal{T}_i not incident to a leaf of

\mathcal{T}_i , and the resulting graph is precisely a graph obtained from $H = (S, E_S)$ by subdividing each edge at most once, and thus H is a minor of G . \square

3. HYPERGRAPHS AND DENSITY

A *partial hypergraph* of \mathcal{H} is a hypergraph obtained from \mathcal{H} by removing a (possibly empty) subset of the edges. In addition to hypergraphs, it will also be convenient to consider *multi-hypergraphs*, i.e. hypergraphs $\mathcal{H} = (V, \mathcal{E})$ where \mathcal{E} is a *multiset* of edges. The *rank* of a hypergraph or multi-hypergraph \mathcal{H} is the maximum cardinality of an edge of \mathcal{H} .

We start with a useful tool, inspired by [15] (see also [4]), itself inspired by the Crossing lemma. Given a graph $G = (V, E)$, we denote by $\text{ad}(G)$ the average degree of G , that is $\text{ad}(G) = 2|E|/|V|$.

Lemma 3. *Let $\mathcal{H} = (V, \mathcal{E})$ be a multi-hypergraph of rank at most $k \geq 2$ on n vertices, and let $E \subseteq \binom{V}{2}$ be a set of pairs of vertices $\{u, v\}$ of V such that there exists an edge e_{uv} of \mathcal{H} containing u and v . (Note that we allow that $e_{uv} = e_{xy}$ for two different pairs $\{u, v\}$ and $\{x, y\}$.) Then the graph (V, E) contains a subgraph H such that $\text{ad}(H) \geq \frac{2|E|}{nek}$ and for every edge uv of H , the corresponding edge e_{uv} of \mathcal{H} contains no vertex from $V(H) - \{u, v\}$.*

Proof. Let \mathbf{H} be the (random) graph obtained by selecting each vertex of \mathcal{H} independently with probability $1/k$, and keeping a single edge (of cardinality 2) between u and v whenever the only selected vertices of e_{uv} are u and v . Then we have

$$\mathbb{E}(|V(\mathbf{H})|) = \frac{n}{k}, \quad \text{and}$$

$$\mathbb{E}(|E(\mathbf{H})|) \geq |E| \cdot \frac{1}{k^2} \left(1 - \frac{1}{k}\right)^{k-2} > \frac{|E|}{ek^2},$$

since $k \geq 2$. It follows that $\mathbb{E}\left(2|E(\mathbf{H})| - \frac{2|E|}{nek}|V(\mathbf{H})|\right) > 0$. In particular, there exists a subgraph H of (V, E) such that $\text{ad}(H) \geq \frac{2|E|}{nek}$ and for every edge uv of H , the edge e_{uv} of \mathcal{H} contains no vertex from $V(H) - \{u, v\}$, as desired. \square

The proof actually gives a randomized algorithm producing the graph H . This algorithm can easily be derandomized using the method of conditional expectations, giving a deterministic algorithm running in time $O(|E| + n)$.

Given a hypergraph \mathcal{H} and a matching \mathcal{B} in \mathcal{H} , we define the *packing-hypergraph* $\mathcal{P}(\mathcal{H}, \mathcal{B})$ as the hypergraph with vertex set \mathcal{B} , in which a subset $\mathcal{B}' \subseteq \mathcal{B}$ is an edge if some edge of \mathcal{H} intersects all the edges in \mathcal{B}' and no other edge of \mathcal{B} .

Lemma 4. *Let G be a graph such that each minor of G has average degree at most d , let \mathcal{H} be a ball hypergraph of G , and let \mathcal{B} be a matching of size n in \mathcal{H} . For every integer $k \geq 2$, the number of edges of cardinality at most k in the packing-hypergraph $\mathcal{P}(\mathcal{H}, \mathcal{B})$ is at most*

$$(1 + dek)^{k-1} \cdot n.$$

Proof. Let \mathcal{P}' be the partial hypergraph of $\mathcal{P}(\mathcal{H}, \mathcal{B})$ induced by the edges of cardinality at most k . Let H be the graph with vertex set \mathcal{B} in which two distinct vertices are adjacent if they are contained in an edge of \mathcal{P}' (i.e. an edge of $\mathcal{P}(\mathcal{H}, \mathcal{B})$ of cardinality at most k). Let m be the number of edges of H . Applying Lemma 3 to \mathcal{P}' , we obtain a subgraph H' of H of average degree at least $\frac{2m}{nek}$, and such that for any pair x, y of adjacent vertices in H' , there is an edge of \mathcal{P}' that contains x and y and no other vertex of H' . The vertices of H' correspond

to a subset S of pairwise disjoint balls of G (since \mathcal{B} is a matching), and each edge of H' corresponds to a ball of G that intersects some pair of balls of S (and does not intersect any other ball of S).

By Lemma 2, H' is a minor of G , so in particular $\frac{2m}{nek} \leq \text{ad}(H') \leq d$, and hence $m \leq \frac{1}{2} dekn$. It follows that H contains a vertex of degree at most dek , and the same is true for every induced subgraph of H (since we can replace \mathcal{B} in the proof by any subset of \mathcal{B}). As a consequence, H is $\lfloor dek \rfloor$ -degenerate. It is a folklore result that ℓ -degenerate graphs on n vertices have at most $\binom{\ell}{t-1} n$ cliques of size t (see for instance [28, Lemma 18], where the proof gives a linear time algorithm to enumerate all the cliques of size t when t and ℓ are fixed), and hence there are at most

$$n \cdot \sum_{i=1}^k \binom{\lfloor dek \rfloor}{i-1} \leq n \cdot (1 + dek)^{k-1}$$

cliques of size at most k in H , which is an upper bound on the number of edges of cardinality at most k in $\mathcal{P}(\mathcal{H}, \mathcal{B})$. \square

Note that the proof gives an $O(n)$ time algorithm enumerating all edges of cardinality at most k in the packing-hypergraph $\mathcal{P}(\mathcal{H}, \mathcal{B})$, when k and d are fixed (note that since H is $\lfloor dek \rfloor$ -degenerate, it contains a linear number of edges, and thus the application of Lemma 3 takes time $O(n)$).

4. FRACTIONAL PACKINGS OF BALLS

We now prove Theorem 4. The proof is inspired by ideas from [22].

Proof of Theorem 4. Let \mathcal{H} be a ball hypergraph of G . Since $\nu^*(\mathcal{H})$ is attained and is a rational number (recall that $\nu^*(\mathcal{H})$ is the solution of a linear program with integer coefficients), there exists a multiset \mathcal{B} of p balls of G , such that every vertex $v \in V(G)$ is contained in at most q balls of \mathcal{B} , and $\nu^*(\mathcal{H}) = p/q$ (see for instance [22], where the same argument is applied to fractional cycle packings). We may assume that q is arbitrarily large (by taking k copies of each ball of \mathcal{B} , with multiplicities, for some arbitrarily large constant k), so in particular we may assume that $q \geq 2$. We may also assume that G contains at least one edge (i.e. $d \geq 1$), otherwise the result clearly holds. Enumerate all the balls in \mathcal{B} as B_1, B_2, \dots, B_p (and recall that since \mathcal{B} is a multiset, some balls B_i and B_j might coincide). We may assume that there is no pair of balls B_i, B_j such that $B_i \subset B_j$ (otherwise we can replace B_j by B_i in \mathcal{B} , and we still have a fractional matching). It follows that the balls of \mathcal{B} are pairwise incomparable (as defined at the beginning of Section 2). For any two intersecting balls B_i and B_j we define x_{ij} as a median vertex of B_i and B_j (also defined at the beginning of Section 2). Recall that it implies in particular that whenever B_i and B_j intersect, $x_{ij} \in B_i \cap B_j$, and if B_i and B_j coincide then x_{ij} is the center of B_i and B_j .

We let \mathcal{G} be the intersection graph of the balls in \mathcal{B} , that is $V(\mathcal{G}) = \mathcal{B}$ and two vertices $B_i, B_j \in \mathcal{B} = V(\mathcal{G})$ with $i \neq j$ are adjacent in \mathcal{G} if and only if $B_i \cap B_j \neq \emptyset$. (In particular, there is an edge linking B_i and B_j when B_i and B_j are two copies of the same ball.) Let m be the number of edges of \mathcal{G} . Let \mathcal{B}^* denote the multi-hypergraph with vertex set \mathcal{B} , where for every vertex of G of the form x_{ij} there is a corresponding edge consisting of the balls in \mathcal{B} that contain x_{ij} . Note that two distinct such vertices could possibly define the same edge, which is why edges in \mathcal{B}^* could have multiplicities greater than 1. The multi-hypergraph \mathcal{B}^* has rank at most q and contains p vertices. Note moreover that the number of pairs of vertices B_i, B_j of \mathcal{B}^* with $i \neq j$ such that there exists an edge of \mathcal{B}^* containing B_i and B_j is precisely m .

By Lemma 3 applied to the multi-hypergraph \mathcal{B}^* , we obtain a graph $H = (S, E_S)$ satisfying the following properties:

- $S \subseteq \mathcal{B}$;
- for each edge $\{B_i, B_j\} \in E_S$, x_{ij} is contained in B_i and B_j but in no other ball from S , and
- $\text{ad}(H) \geq \frac{2m}{peq}$.

We would like to apply Lemma 1 to H but this is not immediately possible, since some balls of S might coincide (recall that \mathcal{B} is a multiset), and therefore the centers of the balls of S might not be pairwise distinct. However, observe that if two balls of S coincide, then by definition the two corresponding vertices of H have degree either 0 or 1 in H (and in the latter case the two vertices are adjacent in H). Indeed, if two balls B_i, B_j of S coincide and B_i is adjacent to B_k in H with $k \neq j$, then the only balls of S containing x_{ik} are B_i and B_k , contradicting the fact that x_{ik} is also in B_j .

Let $S_1 \subseteq S$ be the subset of balls of S having multiplicity 1 in S . Since no ball of \mathcal{B} is a strict subset of another ball of \mathcal{B} , the centers of the balls of S_1 are pairwise distinct. As a consequence of the previous paragraph, if we consider the subgraph H_1 of H induced by S_1 , then $\text{ad}(H) \leq \max(1, \text{ad}(H_1))$.

By Lemma 1 applied to the set of balls S_1 in G , we obtain that H_1 is a minor of G and thus $\text{ad}(H_1) \leq d$. It follows that $\frac{2m}{peq} \leq \text{ad}(H) \leq \max(1, d) \leq d$ (since $d \geq 1$). This implies that the average degree $2m/p$ of \mathcal{G} is at most edq . By the Caro-Wei inequality [5, 27] (or Turán's theorem [25]), it follows that \mathcal{G} contains an independent set of size at least

$$\frac{|V(\mathcal{G})|}{\text{ad}(\mathcal{G}) + 1} \geq \frac{p}{edq + 1} = \frac{\nu^*(\mathcal{H})}{ed + 1/q}.$$

An independent set in \mathcal{G} is precisely a matching in \mathcal{H} , and thus $\nu(\mathcal{H}) \geq \frac{1}{ed+1/q} \cdot \nu^*(\mathcal{H})$ and $\nu^*(\mathcal{H}) \leq (ed + 1/q) \cdot \nu(\mathcal{H})$. Since we can assume that q is arbitrarily large, it follows that $\nu^*(\mathcal{H}) \leq ed \cdot \nu(\mathcal{H})$, as desired.

The rest of the result follows from well known results on the average degree of graphs. On the one hand, an easy consequence of Euler's formula is that planar graphs have average degree at most 6. On the other hand, it was proved by Kostochka [18] and Thomason [24] that every K_t -minor-free graph has average degree $O(t\sqrt{\log t})$. \square

The linear program for ν^* has coefficients in $\{0, 1\}$, and can thus be solved in time $O(n^3)$, since we can assume that the balls have pairwise distinct centers (and so the number of variables and inequalities is linear in the number of vertices). The associated rational coefficients w_e can thus be found in time $O(n^3)$. It is then convenient to define w'_e as the largest $\frac{\ell}{n} \leq w_e$ with $\ell \in \mathbb{N}$. Note that the coefficients (w'_e) still satisfy the inequalities of the linear program for ν^* , and their sum is at least $\nu^* - 1$ since we can assume that there are at most n balls (since their centers are pairwise distinct). There is a small loss on the multiplicative constant (compared to the statement of Theorem 4), but we can now assume that in the proof we have $q \leq n$ and thus $p \leq n^2$ and $m = O(n^3)$. It follows that the application of Lemma 3 can be done in time $O(m) = O(n^3)$, and the construction of a stable set of suitable size in \mathcal{G} can also be done in time $O(m) = O(n^3)$. Therefore, the proof of Theorem 4 gives an $O(n^3)$ time algorithm constructing a matching of size $\Omega(\nu^*(\mathcal{H}))$ in \mathcal{H} .

The *VC-dimension* of a hypergraph \mathcal{H} is the cardinality of a largest subset X of vertices such that for every $X' \subseteq X$, there is an edge e in \mathcal{H} such that $e \cap X = X'$. Bousquet and Thomassé [2] proved the following result.

Theorem 5. *If G has no K_t -minor, then every ball hypergraph \mathcal{H} of G has VC-dimension at most $t - 1$.*

A classical result is that for hypergraphs of bounded VC-dimension, $\tau = O(\tau^* \log \tau^*)$. We will use the following precise bound of Ding, Seymour, and Winkler [8].

Theorem 6. *If a hypergraph \mathcal{H} has VC-dimension at most δ , then*

$$\tau(\mathcal{H}) \leq 2\delta\tau^*(\mathcal{H}) \log(11\tau^*(\mathcal{H})).$$

Combining Theorems 4, 5, and 6, and using that $\nu^*(\mathcal{H}) = \tau^*(\mathcal{H})$, we obtain Theorem 3 as a direct consequence.

As before, the linear program for τ^* has coefficients in $\{0, 1\}$, and can thus be solved in time $O(n^3)$, since we can assume that the balls have pairwise distinct centers (and so the number of variables and inequalities is linear in the number of vertices). The associated rational coefficients w_v can thus be found in time $O(n^3)$. Using algorithmic versions of Theorem 6 (see [17, 21]) and the coefficients (w_v) , a transversal of \mathcal{H} of size $O(\tau^* \log \tau^*) = O(\nu \log \nu)$ can be found by a randomized algorithm sampling $O(\tau^* \log \tau^*)$ vertices according to the distribution given by (w_v) , or a deterministic algorithm running in time $O(n(\tau^{*2} \log \tau^*)^t)$. So the overall complexity of obtaining a transversal of the desired size is $O(n^3)$ (randomized) and $O(n^3 + n(\tau^* \log \tau^*)^t)$ (deterministic). In the remainder of the paper, the result will be used when τ^* is a fixed constant, in which case the complexity of the deterministic algorithm is also $O(n^3)$.

5. LINEAR BOUND

In this section we prove Theorem 2. Recall that by Theorem 3, there is a (monotone) function f_t such that $\tau(\mathcal{H}) \leq f_t(\nu(\mathcal{H}))$ for every ball hypergraph \mathcal{H} of a K_t -minor-free graph. In the proof, we write d_t for the supremum of the average degree of G taken over all graphs G excluding K_t as a minor. Recall that $d_t = O(t\sqrt{\log t})$ [18, 24].

Let $t \geq 1$ be an integer and let $c_t := 2 \cdot (1 + \frac{3}{2}d_t^2 e)^{3d_t/2} \cdot f_t(\frac{3}{2}d_t)$. We will prove that every ball hypergraph \mathcal{H} of a K_t -minor-free graph satisfies $\tau(\mathcal{H}) \leq c_t \cdot \nu(\mathcal{H})$.

Proof of Theorem 2. We prove the result by induction on $k := \nu(\mathcal{H})$. The result clearly holds if $k = 0$ so we may assume that $k \geq 1$. If $k \leq \frac{3}{2}d_t$ then by the definition of f_t we have $\tau(\mathcal{H}) \leq f_t(\frac{3}{2}d_t) \leq c_t \leq c_t \cdot k$, as desired.

Assume now that $k \geq \frac{3}{2}d_t$ and for every ball hypergraph \mathcal{H}' of a K_t -minor-free graph with $\nu(\mathcal{H}') < k$, we have $\tau(\mathcal{H}') \leq c_t \cdot \nu(\mathcal{H}')$. Let G be a K_t -minor-free graph and \mathcal{H} be a ball hypergraph of G with $\nu(\mathcal{H}) = k$. Our goal is to show that $\tau(\mathcal{H}) \leq c_t \cdot k$. Note that we can assume that \mathcal{H} is *minimal*, in the sense that no edge of \mathcal{H} is contained in another edge of \mathcal{H} (otherwise we can remove the larger of the two from \mathcal{H} , this does not change the matching number nor the transversal number).

Consider a maximum matching \mathcal{B} (of cardinality k) in \mathcal{H} . Let \mathcal{E}_1 be the set consisting of all the edges of \mathcal{H} that intersect at most $\frac{3}{2}d_t$ edges of \mathcal{B} . Note that each edge of \mathcal{B} lies in \mathcal{E}_1 , and therefore \mathcal{E}_1 is non-empty. By Lemma 4, the packing-hypergraph $\mathcal{P}(\mathcal{H}, \mathcal{B})$ contains at most $(1 + \frac{3}{2}d_t^2 e)^{3d_t/2} \cdot k$ edges of cardinality at most $\frac{3}{2}d_t$. For each such edge e of $\mathcal{P}(\mathcal{H}, \mathcal{B})$, consider the corresponding subset \mathcal{B}_e of at most $\frac{3}{2}d_t$ edges of \mathcal{B} , and the subset \mathcal{E}_e of edges of \mathcal{H} that intersect each ball of \mathcal{B}_e , and no other ball of \mathcal{B} . Denoting by \mathcal{H}_e the partial hypergraph of \mathcal{H} with edge set \mathcal{E}_e , observe that by the maximality of the matching \mathcal{B} we have $\nu(\mathcal{H}_e) \leq \frac{3}{2}d_t$ (since in \mathcal{B} , replacing the edges of \mathcal{B}_e by a matching of \mathcal{E}_e again gives a matching of \mathcal{H}). It

follows that $\tau(\mathcal{H}_e) \leq f(\frac{3}{2}d_t)$. And thus, if we denote by \mathcal{H}_1 the partial hypergraph of \mathcal{H} with edge set \mathcal{E}_1 , we have

$$\tau(\mathcal{H}_1) \leq (1 + \frac{3}{2}d_t^2 e)^{3d_t/2} \cdot f(\frac{3}{2}d_t) \cdot k = \frac{1}{2}c_t \cdot k.$$

Consider now the subset \mathcal{E}_2 consisting of all the edges of \mathcal{H} that intersect more than $\frac{3}{2}d_t$ edges of \mathcal{B} , and let \mathcal{H}_2 be the partial hypergraph of \mathcal{H} with edge set \mathcal{E}_2 . Note that \mathcal{E}_1 and \mathcal{E}_2 partition the edge set of \mathcal{H} and thus $\tau(\mathcal{H}) \leq \tau(\mathcal{H}_1) + \tau(\mathcal{H}_2)$. Let \mathcal{B}_2 be a maximum matching in \mathcal{H}_2 , and let $\ell = \nu(\mathcal{H}_2) = |\mathcal{B}_2|$. Let H be the (bipartite) intersection graph of the edges of $\mathcal{B} \cup \mathcal{B}_2$, i.e. each vertex of H corresponds to an edge of $\mathcal{B} \cup \mathcal{B}_2$, and two vertices are adjacent if the corresponding edges intersect. (The graph is bipartite because \mathcal{B} and \mathcal{B}_2 are matchings.)

Note that since H is bipartite, for every two distinct edges $\{B, B'\}$ and $\{C, C'\}$ of H , the sets $B \cap B'$ and $C \cap C'$ are disjoint. Moreover, no ball of $\mathcal{B} \cup \mathcal{B}_2$ is a subset of another ball of $\mathcal{B} \cup \mathcal{B}_2$, and thus the balls of $\mathcal{B} \cup \mathcal{B}_2$ are pairwise incomparable (as defined at the beginning of Section 2). So, enumerating the balls in $\mathcal{B} \cup \mathcal{B}_2$ as B_1, B_2, \dots, B_n , we can choose, for each edge $\{B_i, B_j\}$ of H , a median vertex x_{ij} of B_i and B_j (also defined at the beginning of Section 2). Recall that $x_{ij} \in B_i \cap B_j$, and thus it follows from the property above that the only balls of $\mathcal{B} \cup \mathcal{B}_2$ containing x_{ij} are B_i and B_j . By Lemma 1, H is a minor of G and thus has average degree at most d_t . On the other hand, the vertices of H corresponding to the edges of \mathcal{B}_2 have degree at least $\frac{3}{2}d_t$ in H , and thus

$$\frac{3}{2}d_t \cdot \ell \leq \frac{1}{2} \text{ad}(H)(k + \ell) \leq \frac{1}{2}d_t \cdot (k + \ell),$$

where the central term counts the number of edges of H . It follows that $\nu(\mathcal{H}_2) = \ell \leq \frac{k}{2}$, and thus by the induction hypothesis we have $\tau(\mathcal{H}_2) \leq c_t \cdot \nu(\mathcal{H}_2) \leq c_t \cdot \frac{k}{2}$. As a consequence,

$$\tau(\mathcal{H}) \leq \tau(\mathcal{H}_1) + \tau(\mathcal{H}_2) \leq \frac{1}{2}c_t \cdot k + c_t \cdot \frac{k}{2} = c_t \cdot k,$$

which concludes the proof of Theorem 2. \square

The first part of the proof of Theorem 2 uses Theorem 3 when ν (and thus τ^* , by Theorem 4) is bounded by a function of the constant t , and in this case, by the discussion after the proof of Theorem 3, a transversal of the desired size can be found deterministically in time $O(n^3)$.

The second part of the proof of Theorem 2 can be made constructive by performing the following small modification. We observe that we have not quite used the fact that \mathcal{B} is a *maximum* matching of \mathcal{H} , simply that it has the property that, for any edge e in the packing-hypergraph $\mathcal{P}(\mathcal{H}, \mathcal{B})$ of cardinality at most $\frac{3}{2}d_t$, the matching number of \mathcal{H}_e is bounded. As we have explained after Lemma 4, such edges can be enumerated in linear time when t is fixed. We can then compute each $\tau^*(\mathcal{H}_e) = \nu^*(\mathcal{H}_e)$ in time $O(n^3)$ and if this value is more than $e d_t \cdot |e|$, then we can find a matching of size more than $|e| = |\mathcal{B}_e|$ in \mathcal{H}_e in time $O(n^3)$ by Theorem 4, and replace \mathcal{B}_e by this larger matching in \mathcal{B} , thus increasing the size of \mathcal{B} (this can be done at most $\nu(\mathcal{H})$ times). On the other hand, if for all the (linearly many) edges e as above, we have $\tau^*(\mathcal{H}_e) \leq e d_t \cdot |e| = O(d_t^2)$, then by Theorem 3, we can find a transversal of size $O(d_t^2 \log d_t)$ in each hypergraph \mathcal{H}_e in time $O(n^3)$. So overall we find a matching \mathcal{B} that has the desired property, and a transversal of the partial hypergraph of \mathcal{H} with edge set \mathcal{E}_1 of the desired size in time $O(\nu(\mathcal{H}) \cdot n^4)$. Taking the induction step into account (which divides ν by at least 2), we obtain a deterministic algorithm constructing a transversal of size $O(\nu(\mathcal{H}))$ in \mathcal{H} , in time $O(\sum_{i \geq 0} \frac{1}{2^i} \cdot \nu(\mathcal{H}) \cdot n^4) = O(\nu(\mathcal{H}) \cdot n^4)$, when t is a fixed constant.

6. CONCLUSION

The proof of Theorem 2 gives a bound of the order of $\exp(t \log^{3/2} t)$ for the constant c_t . It would be interesting to improve this bound to a polynomial in t .

It is also natural to wonder whether Theorem 2 remains true in a setting broader than proper minor-closed classes. Natural candidates are graphs of bounded maximum degree, graphs excluding a topological minor, k -planar graphs, classes with polynomial growth (meaning that the size of each ball is bounded by a polynomial function of its radius, see e.g. [19]), and classes with strongly sublinear separators (or equivalently, classes with polynomial expansion [11]). We now observe that in all these cases, the associated ball hypergraphs do not satisfy the Erdős-Pósa property, even if all the balls have the same radius. That is, we can find r -ball hypergraphs in these classes with bounded ν and unbounded τ . Our construction shows that this is true even in the seemingly simple case of subgraphs of a grid with all diagonals (i.e. strong products of two paths).

Fix two integers k, ℓ with $k \geq 3$, and ℓ sufficiently large compared to k and divisible by $2\binom{k}{2} - 1$. Given k vertices v_0, v_1, \dots, v_{k-1} , an ℓ -broom with root v_0 and leaves v_1, \dots, v_{k-1} is a tree T of maximum degree 3 with root v_0 and leaves v_1, \dots, v_{k-1} such that

- (1) each leaf is at distance ℓ from the root v_0 ,
- (2) the ball of radius $\ell/2$ centered in v_0 in T is a path (called the *handle* of the broom), and
- (3) the distance between every two vertices of degree 3 in T is sufficiently large compared to k .

We now construct a graph $G_{k,\ell}$ as follows. We start with a set X of k vertices x_1, \dots, x_k , and a path of $\binom{k}{2}$ vertices with vertex set $Y = \{y_{\{i,j\}} \mid 1 \leq i < j \leq k\}$, disjoint from X . We then subdivide each edge of the latter path $\frac{\ell}{2} \frac{1}{\binom{k}{2} - 1} - 1$ times, so that the subdivided path has length $\ell/2$. Finally, for each $1 \leq i \leq k$, we add an ℓ -broom T_i with root x_i and leaves $Y_i = \{y_{\{i,j\}} \mid j \neq i\}$.




FIGURE 1. An embedding of the graph $G_{4,\ell}$ in the 2-dimensional grid with all diagonals (the grid itself is not depicted for the sake of clarity).

We first claim that $G_{k,\ell}$ is a subgraph of the 2-dimensional grid with all diagonals (i.e. the strong products of two paths). To see this, place X on a single column on the left, and Y on another column on the right (in the sequence given by the path), at distance ℓ from the

column of X , then draw each of the brooms in the plane (with crossings allowed). Since the distance between two vertices of degree 3 in a broom is sufficiently large compared to k , we can safely embed each topological crossing in the strong product of two edges (see Figure 1 for an example).

Let $\mathcal{H}_{k,\ell}$ be the ℓ -ball hypergraph of $G_{k,\ell}$ obtained by considering all the balls of radius ℓ in $G_{k,\ell}$. We first observe that $\nu(\mathcal{H}_{k,\ell}) = 1$: this follows from the fact that each ball of radius ℓ centered in a vertex that does not belong to the handle of a broom contains all the vertices of Y , while every two vertices on the handles of two brooms T_i and T_j are at distance at most ℓ from $y_{\{i,j\}}$. Finally, for every two vertices x_i and x_j of X , note that $y_{\{i,j\}}$ is the unique vertex of $G_{k,\ell}$ lying at distance at most ℓ from x_i and x_j , and thus $\tau(\mathcal{H}_{k,\ell}) \geq \frac{k}{2}$. It follows that there is no function f such that $\tau(\mathcal{H}) \leq f(\nu(\mathcal{H}))$ for every ball hypergraph of a subgraph of the strong product of two paths (even when all the balls in the ball hypergraph have the same radius).

Acknowledgments. We thank the two anonymous reviewers for their detailed comments and suggestions.

REFERENCES

- [1] G. Borradaile and E.W. Chambers, *Covering Nearly Surface-Embedded Graphs with a Fixed Number of Balls*, Discrete Comput. Geom. **51** (2014), 979–996.
- [2] N. Bousquet and S. Thomassé, *VC-dimension and Erdős-Pósa property of graphs*, Discrete Math. **338** (2015), 2302–2317.
- [3] P.J. Cameron, *Problems from CGCS Luminy, May 2007*, European J. Combin. **31** (2010), 644–648.
- [4] W. Cames van Batenburg, L. Esperet, and T. Müller, *Coloring Jordan regions and curves*, SIAM J. Discrete Math. **31(3)** (2017), 1670–1684.
- [5] Y. Caro, *New results on the independence number*, Technical report, Tel Aviv University, 1979.
- [6] V. Chepoi, B. Estellon, and G. Naves, *Packing and covering with balls on Busemann surfaces*, Discrete Comput. Geom. **57** (2017), 985–1011.
- [7] V. Chepoi, B. Estellon, and Y. Vaxès, *Covering planar graphs with a fixed number of balls*, Discrete Comput. Geom. **37** (2007), 237–244.
- [8] G. Ding, P. D. Seymour, and P. Winkler, *Bounding the vertex cover number of a hypergraph*, Combinatorica **14** (1994), 23–34.
- [9] G. Ducoffe, M. Habib, and L. Viennot, *Diameter computation on H -minor free graphs and graphs of bounded (distance) VC-dimension*, In Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA), 2020.
- [10] Z. Dvořák, *Constant-factor approximation of the domination number in sparse graphs*, European J. Combin. **34(5)** (2013), 833–840.
- [11] Z. Dvořák and S. Norin, *Strongly sublinear separators and polynomial expansion*, SIAM J. Discrete Math. **30(2)** (2016), 1095–1101.
- [12] P. Erdős and L. Pósa, *On independent circuits contained in a graph*, Canad. J. Math. **17** (1965), 347–352.
- [13] B. Estellon, *Algorithmes de couverture et d’augmentation de graphes sous contraintes de distance*, Thèse de doctorat, Université de Marseille, 2007.
- [14] S. Fiorini, N. Hardy, B. Reed, and A. Vetta, *Approximate min-max relations for odd cycles in planar graphs*, Math. Program. Ser. B **110** (2007), 71–91.
- [15] J. Fox and J. Pach, *Touching strings*, manuscript, 2012.
- [16] M.X. Goemans and D.P. Williamson, *Primal-Dual Approximation Algorithms for Feedback Problems in Planar Graphs*, Combinatorica **17** (1997), 1–23.
- [17] D. Haussler and E. Welzl, *ϵ -nets and simplex range queries*, Discrete Comput. Geom. **2** (1987), 127–151.
- [18] A. Kostochka, *Lower bound of the Hadwiger number of graphs by their average degree*, Combinatorica **4** (1984), 307–316.
- [19] R. Krauthgamer and J. Lee, *The intrinsic dimensionality of graphs*, Combinatorica **27** (2003), 438–447.
- [20] J. Li and M. Parter, *Planar diameter via metric compression*, In Proc. 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC), 2019, 152–163.
- [21] J. Matoušek, *Approximations and Optimal Geometric Divide-and-Conquer*, J. Comput. System Sci. **50(2)** (1995), 203–208.

- [22] B. Reed and F. B. Shepherd, *The Gallai-Younger conjecture for planar graphs*, *Combinatorica* **16** (1996), 555–566.
- [23] B. Reed, N. Robertson, P. Seymour, and R. Thomas, *Packing directed circuits*, *Combinatorica* **16** (1996), 535–554.
- [24] A. Thomason, *An extremal function for contractions of graphs*. *Math. Proc. Cambridge Philos. Soc.* **95** (1984), 261–265.
- [25] P. Turán, *On an extremal problem in graph theory*, *Mat. és Fiz. Lapok* (in Hungarian) **48** (1941), 436–452.
- [26] V. Vapnik and A. Chervonenkis, *On the uniform convergence of relative frequencies of events to their probabilities*, *Theor. Prob. Appl.* **16** (1971), 264–280.
- [27] V. Wei, *A lower bound on the stability number of a simple graph*, Technical report, Bell Labs, 1981.
- [28] D. Wood, *Cliques in graphs excluding a complete graph minor*, *Electron. J. Combin.* **23(3)** (2016), P3.18.

LABORATOIRE G-SCOP, CNRS, UNIV. GRENOBLE ALPES, GRENOBLE, FRANCE.

E-mail address: nicolas.bousquet@grenoble-inp.fr

DÉPARTEMENT D'INFORMATIQUE, UNIVERSITÉ LIBRE DE BRUXELLES, BRUSSELS, BELGIUM.

E-mail address: wcomesva@ulb.ac.be

LABORATOIRE G-SCOP, CNRS, UNIV. GRENOBLE ALPES, GRENOBLE, FRANCE.

E-mail address: louis.esperet@grenoble-inp.fr

DÉPARTEMENT D'INFORMATIQUE, UNIVERSITÉ LIBRE DE BRUXELLES, BRUSSELS, BELGIUM.

E-mail address: gjoret@ulb.ac.be

ALGORITHMS RESEARCH GROUP, UNIVERSITY OF BERGEN, BERGEN, NORWAY.

E-mail address: william.lochet@gmail.com

DÉPARTEMENT DE MATHÉMATIQUE, UNIVERSITÉ LIBRE DE BRUXELLES, BRUSSELS, BELGIUM.

E-mail address: camuller@ulb.ac.be

DÉPARTEMENT D'INFORMATIQUE, UNIVERSITÉ LIBRE DE BRUXELLES, BRUSSELS, BELGIUM, AND LABORATOIRE G-SCOP, CNRS, UNIV. GRENOBLE ALPES, GRENOBLE, FRANCE.

E-mail address: francois.piroto@grenoble-inp.fr