

**HAMBURGER BEITRÄGE
ZUR MATHEMATIK**

Heft 555

**Packing minor-closed families of graphs into
complete graphs**

Silvia Messuti, Hamburg

Vojtěch Rödl, Atlanta

Mathias Schacht, Hamburg

Version August 2015

PACKING MINOR-CLOSED FAMILIES OF GRAPHS INTO COMPLETE GRAPHS

SILVIA MESSUTI, VOJTĚCH RÖDL, AND MATHIAS SCHACHT

ABSTRACT. Motivated by a conjecture of Gyárfás, recently Böttcher, Hladký, Piguet, and Taraz showed that every collection T_1, \dots, T_n of trees on n vertices with $\sum_{i=1}^n e(T_i) \leq \binom{n}{2}$ and with bounded maximum degree, can be packed into the complete graph on $(1 + o(1))n$ vertices. We generalise this result where we relax the restriction of packing families of trees to families of graphs of any given non-trivial minor-closed class of graphs.

1. INTRODUCTION

Given graphs H and F , an F -packing of H is a collection of edge-disjoint subgraphs of H that are isomorphic to F . This definition naturally extends to sequences of graphs. In particular, we say that $\mathcal{F} = (F_1, \dots, F_\ell)$ packs into H if there exist edge-disjoint subgraphs $H_1, \dots, H_\ell \subseteq H$ with H_i isomorphic to F_i for every $i \in [\ell]$.

Gyárfás' tree packing conjecture [7] initiated a lot of research and asserts the following for the case where H is a complete graph and \mathcal{F} is a sequence of trees.

Conjecture 1. *Any sequence $\mathcal{T} = (T_1, \dots, T_n)$ of trees of order $v(T_i) = i$ for $i \in [n]$ packs into K_n .*

The difficulty of this conjecture lies in the fact that it asks for a perfect packing, i.e., a packing where all the edges of K_n are used, since each tree has $e(T_i) = i - 1$ edges and hence $\sum_{i \in [n]} e(T_i) = \binom{n}{2}$. Although some special cases were proven (see, e.g., the surveys [8] and [6], and the references in [4]), this conjecture is still widely open.

Recently Böttcher, Hladký, Piguet, and Taraz [4] showed that a restricted approximate version holds. More precisely, they considered a host graph with slightly more than n vertices and trees with bounded maximum degree, while relaxing the assumption on the number of vertices of each tree.

Theorem 2 (Böttcher, Hladký, Piguet, and Taraz). *For any $\varepsilon > 0$ and any $\Delta \in \mathbb{N}$ there exists $n_0 \in \mathbb{N}$ such that for any $n \geq n_0$ the following holds. If $\mathcal{T} = (T_1, \dots, T_n)$ is a sequence of trees satisfying*

- (1) $\Delta(T_i) \leq \Delta$ and $v(T_i) \leq n$ for every $i \in [n]$, and

The first author was supported through a Doctoral completion scholarship of the University of Hamburg, the second author was supported by NSF grant DMS 1301698, and the third author was supported through the Heisenberg-Programme of the Deutsche Forschungsgemeinschaft (DFG Grant SCHA 1263/4-2).

$$(2) \sum_{i=1}^n e(T_i) \leq \binom{n}{2},$$

then \mathcal{T} packs into $K_{(1+\varepsilon)n}$.

In case $(1 + \varepsilon)n$ is not an integer, we should talk about $K_{\lfloor (1+\varepsilon)n \rfloor}$. However, since we provide asymptotical results, we will omit floors and ceilings here.

The proof of Theorem 2 is based on a randomized embedding strategy, which draws some similarities to the semirandom nibble method (see e.g. [2]). Inspired by the result in [4], we obtained a somewhat simpler proof of Theorem 2, which extends from sequences of trees to sequences of graphs contained in any non-trivial minor-closed class.

Theorem 3. *For any $\varepsilon > 0$, $\Delta \in \mathbb{N}$, and any non-trivial minor-closed family \mathcal{G} there exists $n_0 \in \mathbb{N}$ such that for every $n \geq n_0$ the following holds. If $\mathcal{F} = (F_1, \dots, F_n)$ is a sequence of graphs from \mathcal{G} satisfying*

$$(1) \Delta(F_i) \leq \Delta \text{ and } v(F_i) \leq n \text{ for every } i \in [n], \text{ and}$$

$$(2) \sum_{i=1}^n e(F_i) \leq \binom{n}{2},$$

then \mathcal{F} packs into $K_{(1+\varepsilon)n}$.

In the following we will consider graphs that do not contain isolated vertices. In fact, such vertices can easily be embedded after larger components just by picking any vertex of $K_{(1+\varepsilon)n}$ that has not been used before for the same graph. In the proof we split the graphs F_i into smaller pieces by removing a *small* separator, i.e., a small subset of the vertex set. We discuss these concepts and a generalisation of Theorem 3 in the next section.

2. MAIN TECHNICAL RESULT

We shall establish a generalisation of Theorem 3 for graphs with small separators (see Theorem 7 below). In fact, the Separator Theorem of Alon, Seymour, and Thomas [1] will provide the connection between Theorem 3 and slightly more general Theorem 7.

Theorem 4 (Alon, Seymour, and Thomas). *For every non-trivial minor-closed family of graphs \mathcal{G} there exists $c_{\mathcal{G}} > 0$ such that for every graph $G \in \mathcal{G}$ there exists $U \subseteq V(G)$ with $|U| \leq c_{\mathcal{G}}\sqrt{n}$ such that every component of $G - U$ has order at most $n/2$.*

The graphs we consider in our main result satisfy the following property.

Definition 5. *Given $\delta > 0$ and $s \in \mathbb{N}$, a (δ, s) -separation of a graph $G = (V, E)$ with minimum degree $\delta(G) \geq 1$ is a pair (U, \mathcal{C}) satisfying*

$$i. U \subseteq V, |U| \leq \delta v(G) \text{ and}$$

ii. $\mathcal{C} = G[V \setminus U]$, i.e., the subgraph of G induced on $V \setminus U$, has the property that each component of \mathcal{C} has order at least two and at most s .

We refer to U as the separator, and to \mathcal{C} as the component graph of G .

Note that, for technical reasons that will become clear later (see equation (11)), in ii. we only allow components of size at least 2. Although the removal of a separator could induce components of size one, such a separator U^0 of G may yield at most $\Delta|U^0|$ components of size one, because in our setting we only deal with graphs G of bounded degree $\Delta(G) \leq \Delta$. This allows us to add these “few” vertices to U^0 without enlarging it too much, and ensure that the resulting set U complies with the definition above.

Definition 6. A family \mathcal{G} of graphs with minimum degree at least one is (δ, s) -separable if every $G \in \mathcal{G}$ admits a (δ, s) -separation.

We will deduce Theorem 3 from the following result, in which the condition of \mathcal{G} being minor-closed is replaced by the more general property of being (δ, s) -separable.

Theorem 7. For any $\varepsilon > 0$ and $\Delta \in \mathbb{N}$ there exists $\delta > 0$ such that for every $s \in \mathbb{N}$ and any (δ, s) -separable family \mathcal{G} there exists $n_0 \in \mathbb{N}$ such that for every $n \geq n_0$ the following holds. If $\mathcal{F} = (F_1, \dots, F_n)$ is a sequence of graphs from \mathcal{G} satisfying

- (1) $\Delta(F_i) \leq \Delta$ and $v(F_i) \leq n$ for every $i \in [n]$, and
- (2) $\sum_{i=1}^n e(F_i) \leq \binom{n}{2}$,

then \mathcal{F} packs into $K_{(1+\varepsilon)n}$.

As mentioned above, Theorem 3 easily follows from Theorem 7. First we show that for any non-trivial minor-closed family \mathcal{G} and any $\delta > 0$ there is some s such that \mathcal{G} is (δ, s) -separable. Then we use this fact to deduce Theorem 3.

Given a graph $G \in \mathcal{G}$ of order n with minimum degree $\delta(G) \geq 1$ and maximum degree $\Delta(G) \leq \Delta$, we apply Theorem 4 to all components of G that have some size r_0 with $\frac{n}{2} \leq r_0 \leq n$. Since there are at most two such components, at most two applications of Theorem 4 lead to a separator of size at most $2c_{\mathcal{G}}n^{1/2}$ and a set of components all of which have order less than $n/2$. We then apply Theorem 4 to all components of G that have some size r_1 with $\frac{n}{4} \leq r_1 < \frac{n}{2}$ and obtain another separator of size at most $4c_{\mathcal{G}}\left(\frac{n}{2}\right)^{1/2}$. At this point all components have order less than $n/4$. Again, we apply Theorem 4 to all components of some size r_2 with $\frac{n}{8} \leq r_2 < \frac{n}{4}$, add at most $8c_{\mathcal{G}}\left(\frac{n}{4}\right)^{1/2}$ more vertices to the separator, and so on. After $i > 0$ such iterations we obtain a separator $U^0 \subseteq V(G)$ such that

$$|U^0| \leq 2c_{\mathcal{G}}n^{1/2} + 4c_{\mathcal{G}}\left(\frac{n}{2}\right)^{1/2} + \dots + 2^i c_{\mathcal{G}}\left(\frac{n}{2^{i-1}}\right)^{1/2} < 2c_{\mathcal{G}}n^{1/2} \cdot \frac{\sqrt{2}^i - 1}{\sqrt{2} - 1} < 6c_{\mathcal{G}}n^{1/2} 2^{i/2}$$

and each component of $G - U^0$ has order at most $n/2^i$. For given $\delta > 0$ we can apply this with

$$i = \left\lfloor 2 \log_2 \left(\frac{\delta n^{1/2}}{6c_{\mathcal{G}}(\Delta + 1)} \right) \right\rfloor$$

and obtain a separator U_0 of order at most $\delta n/(\Delta + 1)$, and a set of components all of which have order at most $72c_{\mathcal{G}}^2(\Delta + 1)^2/\delta^2$. Note that some of the components in $G - U^0$ may have size one. However, owing to the maximum degree of G there are at most $\Delta|U^0|$ such components. By defining U as the separator of size at most δn obtained from U^0 by adding all these degenerate components of order one, we have shown that the non-trivial minor-closed family \mathcal{G} is (δ, s) -separable with $s = \frac{72c_{\mathcal{G}}^2(\Delta+1)^2}{\delta^2}$. Applying Theorem 7 with this s yields Theorem 3.

The rest of this paper is devoted to the proof of Theorem 7. In Section 3 we introduce some definitions and state two technical lemmas that are used in the proof of the theorem, which is given at the end of the section. Resolvable and almost resolvable decompositions, which we will use to construct our packing, are introduced in Section 4. Finally, the two technical lemmas, Lemma 10 and Lemma 11, are proved in Sections 5 and 6, respectively.

3. PROOF OF THE MAIN RESULT

The following notation will be convenient.

Definition 8. *Let \mathcal{G} be a family of graphs. An n -tuple of graphs $\mathcal{F} = (F_1, \dots, F_n)$ with $F_i \in \mathcal{G}$ and $i \in [n]$ is called a (\mathcal{G}, n, Δ) -sequence if*

- (1) $\Delta(F_i) \leq \Delta$ and $v(F_i) \leq n$ for every $i \in [n]$, and
- (2) $\sum_{i=1}^n e(F_i) \leq \binom{n}{2}$.

In this paper we will consider (\mathcal{G}, n, Δ) -sequences with the following additional properties:

- \mathcal{G} will be a (δ, s) -separable family, and
- each graph F_i will be associated with a fixed (δ, s) -separation (U_i, \mathcal{C}_i) .

For a simpler notation we will often suppress the dependence on U_i when we refer to a (\mathcal{G}, n, Δ) -sequence (F_1, \dots, F_n) , since the separator U_i will be always clear from the context. In a component C from \mathcal{C}_i we distinguish the set of vertices that are connected to the separator U_i and refer to this set as the *boundary* ∂C of C

$$\partial C = V(C) \cap N(U_i),$$

where as usual $N(U_i)$ denotes the union of the neighbours in F_i of the vertices in U_i . Moreover, for a component graph \mathcal{C}_i we consider the union of the boundary sets of its components and set

$$\partial \mathcal{C}_i = \bigcup \{ \partial C : C \text{ component in } \mathcal{C}_i \}.$$

Note that

$$|\partial \mathcal{C}_i| \leq \sum_{u \in U_i} d(u) \leq |U_i| \Delta \leq \delta n \Delta \tag{1}$$

For the proof of Theorem 7 we shall pack a given (\mathcal{G}, n, Δ) -sequence into $K_{(1+\varepsilon)n}$. The vertices of the host graph $K_{(1+\varepsilon)n}$ will be split into a large part X of order $(1 + \xi)n$ for some carefully chosen $\xi = \xi(\varepsilon, \Delta) > 0$, and a small part $Y = V \setminus X$. We will pack the graphs $\{\mathcal{C}_i\}_{i \in [n]}$ into the clique spanned on X and the sets $\{U_i\}_{i \in [n]}$ into Y . For this, we shall ensure that the vertices representing the boundary $\partial\mathcal{C}_i$ will be appropriately connected to the vertices representing the separator U_i . Having this in mind we will make sure that each vertex of X will only host a few boundary vertices. In fact, since every edge of the complete bipartite graph induced by X and Y can be used only once in the packing, each vertex $x \in X$ can be used at most $|Y|$ times as boundary vertex for the packing of the sequence $\{\mathcal{C}_i\}_{i \in [n]}$. This leads to the following definition.

Definition 9. For $i \in [n]$, let $F_i = (V_i, E_i)$ be graphs with separators $U_i, \subseteq V_i$ and component graphs $\mathcal{C}_i = F_i[V_i \setminus U_i]$. For a family of injective maps $\mathbf{f} = \{f_i\}_{i \in [n]}$ with $f_i: V(\mathcal{C}_i) \rightarrow X$ and for $x \in X$ we define the boundary degree of x with respect to \mathbf{f} by

$$d_{\mathbf{f}}^{\partial}(x) = |\{i \in [n]: f_i^{-1}(x) \in \partial\mathcal{C}_i\}|.$$

We call such a family of maps b -balanced for some $b \in \mathbb{R}$ if $d_{\mathbf{f}}^{\partial}(x) \leq b$ for every $x \in X$.

The proof of Theorem 7 will follow from Lemma 10 and Lemma 11 below. Lemma 10 yields a balanced packing of the component graphs $\{\mathcal{C}_i\}_{i \in [n]}$ into the clique spanned by X with $|X| \leq (1 + \xi)n$.

Lemma 10. For any $\xi > 0$ and $\Delta \in \mathbb{N}$ there exists $\delta > 0$ such that for every $s \in \mathbb{N}$ and any (δ, s) -separable family \mathcal{G} there exists $n_0 \in \mathbb{N}$ such that if \mathcal{F} is a (\mathcal{G}, n, Δ) -sequence with $n \geq n_0$, then there exists a (ξn) -balanced packing of the component graphs $\{\mathcal{C}_i\}_{i \in [n]}$ of all members of \mathcal{F} into $K_{(1+\xi)n}$.

Once we have a balanced packing of $\{\mathcal{C}_i\}_{i \in [n]}$ into $K_{(1+\xi)n}$, the next lemma allows us to extend it to a packing of $\mathcal{F} = (F_1, \dots, F_n)$ into a slightly larger clique of size $(1 + \varepsilon)n$.

Lemma 11. For any $\varepsilon > 0$ and $\Delta \in \mathbb{N}$, there exist $\xi > 0$ and $\delta > 0$ such that for every s and any (δ, s) -separable family \mathcal{G} there exists n_0 such that for any $n \geq n_0$ the following holds. Suppose there exists a (ξn) -balanced packing of the component graphs $\{\mathcal{C}_i\}_{i \in [n]}$ associated with a (\mathcal{G}, n, Δ) -sequence \mathcal{F} into $K_{(1+\xi)n}$. Then there exists a packing of \mathcal{F} into $K_{(1+\varepsilon)n}$.

We postpone the proofs of Lemma 10 and Lemma 11 to Section 5 and Section 6, respectively. Here we describe the proof of our main Theorem based on these two lemmas.

3.1. Proof of Theorem 7. We will first fix all involved constants. Note that Theorem 7 and Lemma 11 have a similar quantification. Hence, for the proof of Theorem 7, we may apply Lemma 11 with ε and Δ from Theorem 7 and obtain ξ and δ' . Then Lemma 10 applied

with ξ and Δ yields a constant δ'' . For Theorem 7 we set $\delta = \min\{\delta', \delta''\}$. After displaying δ for Theorem 7 we are given some $s \in \mathbb{N}$ and a (δ, s) -separable family \mathcal{G} .

With constants chosen as above, we can apply Lemma 10 for a (\mathcal{G}, n, Δ) -sequence \mathcal{F} which then asserts that the assumptions of Lemma 11 are fulfilled. Finally, the conclusion of Lemma 11 yields Theorem 7. \square

4. RESOLVABLE AND ALMOST RESOLVABLE DECOMPOSITIONS

The idea of the proof is to split the given graphs (F_1, \dots, F_n) into small components, group such components by isomorphism types, and pack components from the same group into complete subgraphs of $K_{(1+\varepsilon)n}$. For that we will use Theorem 12 and Theorem 14.

A K_m -factor of K_n is a collection of $\frac{n}{m}$ vertex disjoint cliques of order m , and a resolvable K_m -decomposition of K_n is a collection of

$$\frac{\binom{n}{2}}{\binom{m}{2}} \frac{m}{n} = \frac{n-1}{m-1}$$

edge disjoint K_m -factors. Theorem 12 states that the obvious necessary divisibility conditions for the existence of a K_m -decomposition of K_n are actually sufficient.

Theorem 12 (Ray-Chaudhury and Wilson). *For every $m \geq 2$ there exists n_0 such that if $n \geq n_0$ and $n \equiv m \pmod{m(m-1)}$, then K_n admits a resolvable K_m -decomposition.*

For general F , resolvable decompositions do not necessarily exist (for example it is easy to see that there is no n for which resolvable $K_{1,3}$ -decompositions of K_n exist). Therefore, instead of F -factors, we consider F -matchings, i.e., sets of vertex disjoint copies of F .

Definition 13. *An (F, η) -factorization of K_ℓ is a collection of F -matchings of K_ℓ such that*

- i. each matching has size at least $(1 - \eta) \frac{\ell}{v(F)}$, and*
- ii. these matchings together cover all but at most $\eta \binom{\ell}{2}$ edges of K_ℓ .*

From these two properties we deduce that the number t of F -matchings in an (F, η) -factorization satisfies

$$(1 - \eta) \frac{(\ell - 1)v(F)}{2e(F)} \leq t \leq \frac{(\ell - 1)v(F)}{2e(F)}.$$

Also note that any $(F, 0)$ -factorization of K_ℓ is a resolvable F -decomposition of K_ℓ . We will then use the following approximate result, which can be deduced from [5] and [10] (see also [3]).

Theorem 14. *For every F and $\eta > 0$ there exists ℓ_0 such that for every $\ell \geq \ell_0$ there exists an (F, η) -factorization of K_ℓ .*

5. PACKING THE COMPONENTS

The crucial part in the proof of Theorem 7 is Lemma 10, which we are going to prove in this section. In Lemma 10 we are given a (\mathcal{G}, n, Δ) -sequence (F_1, \dots, F_n) of graphs from a (δ, s) -separable family \mathcal{G} with fixed separations (U_i, \mathcal{C}_i) associated with each F_i . Our goal will be to construct a (ξn) -balanced packing of the component graphs $\{\mathcal{C}_i\}_{i \in [n]}$ into K_N , with $N = (1 + \xi)n$.

The packing of $\{\mathcal{C}_i\}_{i \in [n]}$ will make use of a resolvable K_m -decomposition of K_N (actually we will use a somewhat more complicated auxiliary structure which we will describe in Section 5.1) and will be realized in two steps: the *assignment* phase and the *balancing* phase.

- In the assignment phase we consider a K_m -decomposition of K_N and then describe which components of each \mathcal{C}_i are assigned to which copies of K_m .
- In the balancing phase we make sure that the mapping from components of each \mathcal{C}_i into copies of K_m from K_N will form a (ξn) -balanced packing as promised in Lemma 10.

Below we outline the main ideas of these two steps. We start with the assignment phase first. The balancing phase will be discussed in Section 5.3.

5.1. Outline of the assignment phase. The purpose of the assignment phase is to produce a “preliminary packing” of each \mathcal{C}_i , $i = 1, \dots, n$ into some K_m -factor. We recall that each component graph \mathcal{C}_i consists of several components each with at most s vertices and maximum degree at most Δ . Moreover, in each component C we distinguish the set ∂C of vertices that are connected to the separator U_i .

We define an *isomorphism type* S as a pair (R, B) where R is a graph on at most s labeled vertices and maximum degree at most Δ , and B is a subset of the vertices of R . Let $\mathcal{S} = (S_1, \dots, S_\sigma)$ be the enumeration of all isomorphism types $S_j = (R_j, B_j)$, such that

$$\frac{e(R_1)}{v(R_1)} \geq \dots \geq \frac{e(R_\sigma)}{v(R_\sigma)}. \quad (2)$$

The definition of \mathcal{S} yields

$$\sigma \leq 2^{\binom{s}{2}} \cdot 2^s \leq 2^{s^2}. \quad (3)$$

Note for every component C of \mathcal{C}_i there exists an isomorphism type $S_j = (R_j, B_j) \in \mathcal{S}$ such that there exists a graph isomorphism $\phi: V(C) \rightarrow V(R_j)$ with the additional property that $\phi(\partial C) = B_j$. Therefore, we can describe the structure of a component graph \mathcal{C}_i as a disjoint union

$$\mathcal{C}_i = \bigcup_{S \in \mathcal{S}} \nu_i(S) \cdot S$$

where $\nu_i(S)$ denotes the number of components isomorphic to S contained in \mathcal{C}_i . In the rest of the paper we will simplify the notation and refer to S as a graph.

The assignment procedure makes use of further decomposition layers. In fact, for each copy of K_m appearing in the resolvable decomposition of K_N we consider a resolvable K_ℓ -decomposition of such a copy of K_m . Each resolution class consisting of $\frac{m}{\ell}$ disjoint copies of K_ℓ will be reserved for some isomorphism class S and the copies of S coming from various \mathcal{C}_i will be then packed into each such K_ℓ . Since we consider K_m -decomposition of K_N , K_ℓ -decomposition of K_m , and S -decomposition of K_ℓ for each $S \in \mathcal{S}$, we will refer to such structure as *three layer decomposition* and motivate its use below.

5.2. The three layer decomposition. We begin our discussion with the simpler case when all components in all the component graphs \mathcal{C}_i are isomorphic to a given graph S and argue why even in this simpler case at least two layers are required. Then we look at the general case, where the component graphs consist of more different isomorphism types, and explain the use of three layers.

5.2.1. One layer. In the case where all components in $\{\mathcal{C}_i\}_{i \in [n]}$ are isomorphic to a single graph S , a straightforward way to pack $\{\mathcal{C}_i\}_{i \in [n]}$ into K_N would be the following. Suppose there exists a resolvable S -decomposition of K_N . Then, by assigning the components of a graph \mathcal{C}_i to copies of S from the same S -factor, we ensure that the components within each component graph are packed vertex-disjointly.

With this approach, however, we might end up not covering many edges of K_N (and consequently not being able to find a packing of the graphs \mathcal{C}_i). Let \mathcal{C}_1 and \mathcal{C}_2 be component graphs with strictly more than $N/2$ vertices. Once we assign the components of \mathcal{C}_1 to an S -factor of K_N , we cannot use the other copies of S in the same S -factor to accommodate the components of \mathcal{C}_2 . In fact, at least one component of \mathcal{C}_2 would not fit in that S -factor and we would have to use a copy of S from another S -factor. We would have to ensure that this copy of S is vertex disjoint from those already used for \mathcal{C}_2 in the previous S -factor, and an obvious way to get around this would be to embed all components of \mathcal{C}_2 in a new S -factor all together. However, this would be very wasteful and if many (for example $\Omega(n)$) graphs \mathcal{C}_i would be of size strictly larger than $N/2$, then we would not be able to pack all \mathcal{C}_i into K_N in such a straightforward way. We remedy this situation by introducing an additional layer.

5.2.2. Two layers. For an appropriately chosen integer m , suppose there exist a resolvable K_m -decomposition of K_N and a resolvable S -decomposition of K_m . Note that, with this additional decomposition layer at hand, we can address the issue raised above more easily. In fact, we fix a K_m -factor of K_N and use sufficiently many K_m 's of this K_m -factor to host the components of \mathcal{C}_1 , all of which are isomorphic to S by our assumption. The remaining K_m 's of the factor can host the first part of \mathcal{C}_2 . We then “wrap around” and reuse the K_m 's

containing copies of S from \mathcal{C}_1 by selecting a new S -factor inside these K_m 's to host the second part of \mathcal{C}_2 . This way the components of \mathcal{C}_1 and \mathcal{C}_2 are packed edge disjointly and the components of \mathcal{C}_2 (resp. \mathcal{C}_1) are in addition vertex disjoint, as required for a packing. We can continue this process to pack $\mathcal{C}_3, \mathcal{C}_4, \dots$ until the fixed K_m -factor of K_N is fully used. Then we continue with another K_m -factor and so on.

This procedure will work if all components of each \mathcal{C}_i are isomorphic to a single S . Let us note however that in case \mathcal{C}_i contains components of different isomorphism types two layers may not be sufficient. This is because we would have to select S -factors for different graphs S within K_m and there seems to be no obvious way to achieve this in a two layer decomposition. Instead we will introduce a third layer, which will give us sufficient flexibility to address this issue.

5.2.3. *Three layers.* Here we give an outline describing how a three layer structure can be used to address the general problem. The details will follow in section 5.4.1. Consider a resolvable K_m -decomposition $\mathcal{D}^{m,N}$ of K_N , a resolvable K_ℓ -decomposition $\mathcal{D}^{\ell,m}$ of K_m , and resolvable S -decompositions $\mathcal{D}^{S,\ell}$ of K_ℓ for every $S \in \mathcal{S}$ (in fact the last assumption will never be used in its full strength, we will use Theorem 14 instead). We view resolvable decompositions as collections of factors. In particular we write, e.g., $\mathcal{D}^{m,N} = \{\mathcal{D}_1^{m,N}, \dots, \mathcal{D}_{\frac{N-1}{m-1}}^{m,N}\}$, where $\mathcal{D}_k^{m,N}$ is a K_m -factor of K_N for $k = 1, \dots, \frac{N-1}{m-1}$.

Suppose now we are given graphs $\mathcal{C}_1, \dots, \mathcal{C}_n$, $\mathcal{C}_i = \bigcup_{S \in \mathcal{S}} \nu_i(S) \cdot S$. We will proceed greedily processing the \mathcal{C}_i 's one by one. In each step we will work with one fixed K_m -factor $\mathcal{D}_k^{m,N} = \mathcal{D}_{\text{current}}^{m,N}$ of K_N which will be used repeatedly as long as ‘‘sufficiently many’’ edges of such factor are available. For example, $\mathcal{D}_1^{m,N}$ will host $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_a$ for some $a < n$, then $\mathcal{D}_2^{m,N}$ will host $\mathcal{C}_{a+1}, \mathcal{C}_{a+2}, \dots, \mathcal{C}_b$ for some $a < b < n$, and so on. Once we run out of available edges in factor $\mathcal{D}_{\text{current}}^{m,N}$ we will move that factor in the set $\mathcal{D}_{\text{used}}^{m,N} \subseteq \mathcal{D}^{m,N}$ of factors the edges of which were already assigned to previous \mathcal{C}_i 's and select a new factor $\mathcal{D}_{\text{current}}^{m,N} \in \mathcal{D}^{m,N} \setminus \mathcal{D}_{\text{used}}^{m,N}$ which we will continue to work with.

Let us now outline the assignment within a K_m of the current K_m -factor. For each $K_m \in \mathcal{D}_{\text{current}}^{m,N}$ we consider a resolvable decomposition $\mathcal{D}^{\ell,m} = \mathcal{D}^{\ell,m}(K_m)$ of such a K_m . Again some factors in that decomposition might have already been completely used. Among those which were not completely used yet, we specify σ of such ‘‘current’’ factors $\mathcal{D}_S^{\ell,m}$, each ready to be used to embed copies of S in the current particular step. Since K_ℓ admits resolvable S -decompositions for every $S \in \mathcal{S}$, each $\mathcal{D}_S^{\ell,m}$ corresponds to $\frac{(\ell-1)v(S)}{2e(S)} = t(S)$ S -factors of K_m which we may denote by $\mathcal{D}_1^{S,\ell,m}, \dots, \mathcal{D}_{t(S)}^{S,\ell,m}$. At each step, in every K_m we will only use one of such S -factors, which we denote by $\mathcal{D}_{\text{current}}^{S,\ell,m}$. A set of components of \mathcal{C}_i that are going to be assigned to an S -factor of a K_m will be referred to as a *chunk*.

With this structure in mind we are able to describe our greedy assignment procedure. Assume that in the assignment procedure the graphs $\mathcal{C}_1, \dots, \mathcal{C}_{i-1}$ were already processed and that $\mathcal{C}_i = \bigcup_{S \in \mathcal{S}} \nu_i(S) \cdot S$. The assignment of \mathcal{C}_i will consist of the following four steps which we discuss in detail in Section 5.4.1.

- i. For every isomorphism type $S \in \mathcal{S}$, partition the $\nu_i(S)$ components into as few as possible chunks of size at most $\frac{m}{v(S)}$.
- ii. For every $S \in \mathcal{S}$, select $\frac{\nu_i(S)v(S)}{m}$ copies of K_m from the current K_m -factor $\mathcal{D}_{\text{current}}^{m,N}$ and match each such K_m with a chunk of components isomorphic to S .
- iii. For every $S \in \mathcal{S}$ and for each chunk of type S , assign the components in the chunk to the S -factor $\mathcal{D}_{\text{current}}^{S,\ell,m}$ of K_m . The copies of S will cover $m \frac{e(S)}{v(S)}$ edges of K_m .
- iv. Prepare for the assignment of the next component graph.

This procedure leads to a packing of $\{\mathcal{C}_i\}_{i \in [n]}$ into K_N if we do not run out of K_m -factors during the process, and in the proof we shall verify this. Assuming this for the moment, the procedure above yields a preliminary packing which can be encoded by functions $\mathbf{f} = \{f_i\}_{i \in [n]}$, with $f_i: V(\mathcal{C}_i) \rightarrow V(K_N)$.

5.3. Outline of the balancing phase. In this section we will outline how the preliminary packing \mathbf{f} obtained in the assignment phase is used to realize a (ξn) -balanced packing of $\{\mathcal{C}_i\}_{i \in [n]}$ into K_N . Further detail will be given in Section 5.4.2.

Note that so far we did not consider the boundary degrees of the vertices of K_N and, in fact, \mathbf{f} is not guaranteed to be balanced. However, the layered structure of the assignment will allow us to fix this by using the following degrees of freedom. Firstly, the $\frac{N}{m}$ K_m 's in any of the $\frac{N-1}{m-1}$ K_m -factors from $\mathcal{D}^{m,N}$ can be permuted independently for each K_m -factor. Since any component graph is assigned to a single K_m -factor, the resulting mappings remain injective and the embedding of the \mathcal{C}_i 's stays pairwise edge disjoint. Secondly, each K_m can be embedded into K_N in $m!$ possible ways by permuting its vertices. There are

$$\left(\left(\frac{N}{m} \right)! \times (m!)^{\frac{N}{m}} \right)^{\frac{N-1}{m-1}}$$

such choices in total and each of them leads to a packing of the component graphs $\{\mathcal{C}_i\}_{i \in [n]}$.

We will pick one of such choices uniformly at random, and show that with positive probability each vertex of K_N is used as a boundary vertex approximately the same number of times. Since the sum of the boundary degrees is at most $\Delta \delta n^2 \leq \xi n^2 / 2$ (see (1)), this leads to a (ξn) -balanced packing \mathbf{g} of $\{\mathcal{C}_i\}_{i \in [n]}$ into K_N .

5.4. Proof of Lemma 10. Given ξ and Δ , set

$$\delta = \xi / 2\Delta \tag{4}$$

and let \mathcal{G} be a (δ, s) -separable family, for some $s \in \mathbb{N}$. We apply Theorem 14 with

$$\eta = \xi/8 \tag{5}$$

and fix an integer $\ell > s^2$ satisfying that for every $S \in \mathcal{S}$ there exists an (S, η) -factorization of K_ℓ . Let $m \in \mathbb{N}$ such that

$$m > 16\sigma\ell/\xi \tag{6}$$

and there exists a resolvable K_ℓ -decomposition of K_m (see Theorem 12). Similarly, let

$$n_0 > \max\{4m^2/\xi, 2^{2m}\} \tag{7}$$

such that for any $n \geq n_0$ satisfying the necessary congruence property there exists a resolvable K_m -decomposition of K_n . Having defined n_0 , we are now given a (\mathcal{G}, n, Δ) -sequence $\mathcal{F} = (F_1, \dots, F_n)$ for some $n \geq n_0$. We will show that there exists a (ξn) -balanced packing of the family of component graphs $\{\mathcal{C}_i\}_{i \in [n]}$ into K_N , for any N with $(1 + \frac{\xi}{2})n \leq N \leq (1 + \xi)n$ such that K_N admits a K_m -decomposition. Since $n \geq n_0 \geq \frac{4m^2}{\xi}$, such N indeed exist.

5.4.1. *The assignment phase.* Next we elaborate on the outline given in Sections 5.1 and 5.2. First we describe the auxiliary structure we are going to use followed by the actual assignment procedure.

The auxiliary structure. For each $S \in \mathcal{S}$ let $\mathcal{D}^{S, \ell}$ be a fixed (S, η) -factorization of K_ℓ (see Definition 13). Let $\mathcal{D}^{\ell, m}$ be an arbitrarily chosen resolvable K_ℓ -decomposition of K_m . Similarly, for the given N , denote by $\mathcal{D}^{m, N}$ an arbitrarily chosen resolvable K_m -decomposition of K_N .

At each point of time in the assignment procedure we will work with one K_m -factor which we refer to as the *current K_m -factor* $\mathcal{D}_{\text{current}}^{m, N} \in \mathcal{D}^{m, N}$. Each K_m of the current K_m -factor is decomposed into K_ℓ -factors using $\mathcal{D}^{\ell, m}$. Moreover, in every K_m , for every $S \in \mathcal{S}$ we pick a K_ℓ -factor which we denote by $\mathcal{D}_S^{\ell, m}$. We refer to $\mathcal{D}_S^{\ell, m}$ as the *current K_ℓ -factor for S* . We then apply Theorem 14 to all K_ℓ 's in such a K_ℓ -factor and obtain (S, η) -factorizations for every K_ℓ in $\mathcal{D}_S^{\ell, m}$. Note that we can arbitrarily fix an S -matching in each K_ℓ of $\mathcal{D}_S^{\ell, m}$ and obtain an S -matching of K_m of size at least

$$(1 - \eta) \frac{\ell}{v(S)} \frac{m}{\ell} = (1 - \eta) \frac{m}{v(S)}. \tag{8}$$

This way we set up $t(S)$ edge disjoint S -matchings of K_m contained in $\mathcal{D}_S^{\ell, m}$, for

$$(1 - \eta) \frac{(\ell - 1)v(S)}{2e(S)} \leq t(S) \leq \frac{(\ell - 1)v(S)}{2e(S)},$$

which we denote by $\mathcal{D}_1^{S, \ell, m}, \dots, \mathcal{D}_{t(S)}^{S, \ell, m}$. Note that such S -matchings cover at least $(1 - \eta) \frac{m}{\ell} \binom{\ell}{2}$ edges of the K_ℓ 's in $\mathcal{D}_S^{\ell, m}$.

Every such structure will be used until it is considered *full* according to the following definition.

Definition 15. A K_ℓ -factor $\mathcal{D}_S^{\ell,m}$ is full when all its S -matchings have been used. A K_m is full when there exists an isomorphism type $S \in \mathcal{S}$ such that $\mathcal{D}_S^{\ell,m}$ is full and any other K_ℓ -factor is either full or reserved to another isomorphism type. A K_m -factor is full when one of its K_m 's is full.

The assignment procedure. We now give the details of the four steps outlined in Section 5.2.3 for the assignment for the graph $\mathcal{C}_i = \bigcup_{S \in \mathcal{S}} \nu_i(S) \cdot S$. We assume that the graphs $\mathcal{C}_1, \dots, \mathcal{C}_{i-1}$ have already been assigned and that the current K_m -factor $\mathcal{D}_{\text{current}}^{m,N} = \mathcal{D}_k^{m,N}$ is not full.

- i) For each isomorphism type $S \in \mathcal{S}$ we group the $\nu_i(S)$ copies of S into as few as possible chunks of size at most $(1 - \eta) \frac{m}{v(S)}$ (note that this matches the size of an S -matching of K_m , as given in (8)). The correction factor $(1 - \eta)$ here addresses the fact that we deal with (S, η) -factorizations and not with resolvable S -decompositions. The number $\mu_i(S)$ of chunks required for the $\nu_i(S)$ components of type S is hence given by

$$\mu_i(S) = \left\lceil \frac{\nu_i(S) \cdot v(S)}{(1 - \eta)m} \right\rceil. \quad (9)$$

- ii) We order the K_m 's in the current K_m -factor $\mathcal{D}_{\text{current}}^{m,N}$ according to the number of edges that have already been assigned to it. We start with the one in which the least number of edges have been used. We then assign the $\mu_i(S_1)$ chunks of type S_1 to the first $\mu_i(S_1)$ K_m 's in that order and continue in the natural way, that is, the $\mu_i(S_2)$ chunks of type S_2 are assigned to the next $\mu_i(S_2)$ K_m 's, and so on. Since the members of \mathcal{S} are ordered non-increasingly according to their densities (see 2), this way we will ensure that the K_m 's in the current K_m -factor are used in a balanced way, which is essential to leave only little waste.
- iii) Once we have determined which chunk goes to which K_m , we have to assign the components S of the chunk to their copies in the corresponding K_m . In the chosen K_m we assign the components of the chunk to $\mathcal{D}_{\text{current}}^{S,\ell,m}$. Such a matching exists because we assumed that the current K_m -factor $\mathcal{D}_k^{m,N}$ is not full. Note that, independently of the precise number of components in the chunk, we use an entire S -matching in all the K_ℓ 's of the current K_ℓ -factor for S for the assignment of this chunk.
- iv) After we have assigned the components of \mathcal{C}_i we prepare for the assignment of \mathcal{C}_{i+1} . In each K_m , for every isomorphism type S , we check whether the current K_ℓ -factor for S is full. If it is, two cases may arise. In the first case there exists another K_ℓ -factor in the K_m that has not been reserved for any $S \in \mathcal{S}$ yet. Then, we apply Theorem 14

with S and η to all copies of K_ℓ in such a K_ℓ -factor and this factor becomes the current K_ℓ -factor for S , i.e., $\mathcal{D}_S^{\ell,m}$ in that K_m . In the second case, all K_ℓ -factors are either full or have been reserved for some $S' \in \mathcal{S}$ with $S' \neq S$, hence we cannot set up a new K_ℓ -factor for S . This implies that the K_m and the K_m -factor are full (see Definition 15). Since we assigned the components of \mathcal{C}_i to the least used K_m 's in the K_m -factor, we are ensured that at this point all the K_m 's in $\mathcal{D}_{\text{current}}^{m,N}$ are almost completely used. At this point we add $\mathcal{D}_{\text{current}}^{m,N}$ to $\mathcal{D}_{\text{used}}^{m,N}$ and set $\mathcal{D}_{\text{current}}^{m,N} = \mathcal{D}_{k+1}^{m,N}$.

The assignment phase yields a packing. We shall verify that the procedure yields a correct assignment. For that we have to show that any component graph \mathcal{C}_i “fits” into K_N , and that we do not run out of K_m -factors while iterating the four steps for all graphs in $\{\mathcal{C}_i\}_{i \in [n]}$.

We first show that every \mathcal{C}_i fits into one K_m -factor. Recall that in Step i) the copies isomorphic to some $S \in \mathcal{S}$ are split into chunks of size at most $(1 - \eta)\frac{m}{v(S)}$ and each chunk is assigned to an S -matching of $\mathcal{D}_S^{\ell,m}$. At this point some vertices may not be used for one of the following two reasons:

- V1) We always reserve a whole S -matching $\mathcal{D}_{\text{current}}^{S,\ell,m}$ for each chunk, even though some chunks may contain only a few copies of S . In the worst case where only one copy of S is contained in the chunk we may waste $m - v(S) \leq m$ vertices and in principle this could happen for every isomorphism type $S \in \mathcal{S}$. However, since such a “rounding error” occurs at most once for each isomorphism type, we may waste at most σm vertices for this reason.
- V2) We cannot guarantee that the S -matchings which we are using are perfect S -factors. However, from Theorem 14 it follows that each matching covers at least $(1 - \eta)\frac{m}{v(S)}v(S) = (1 - \eta)m$ vertices of K_m . Therefore the number of uncovered vertices in the K_m -factor due to this imperfection is at most $\eta m \frac{N}{m} = \eta N$.

Hence \mathcal{C}_i fits into one K_m -factor if we ensure that $v(\mathcal{C}_i) + \sigma m + \eta N \leq N$, which follows from

$$v(\mathcal{C}_i) + \sigma m + \eta N \leq n + \sigma m + \eta N \leq (1 + \frac{\xi}{2})n \leq N,$$

due to (5), (6), and (7).

It is left to show that $\frac{N-1}{m-1}$ K_m -factors are sufficient to host all the graphs from $\{\mathcal{C}_i\}_{i \in [n]}$. For that, we shall bound the number of unused edges in each K_m -factor. At the point when a K_m becomes full, all its K_ℓ -factors, except for the current K_ℓ -factors $\mathcal{D}_S^{\ell,m}$ for each isomorphism type $S \in \mathcal{S}$, have been used in the assignment. This leads to the following cases.

- E1) The current K_ℓ -factor $\mathcal{D}_S^{\ell,m}$ for a given isomorphism type S may not have been used at all and hence all its $\binom{\ell}{2} \frac{m}{\ell}$ edges are not used in the assignment.
- E2) Owing to Theorem 14, in a used K_ℓ -factor, up to at most $\eta \binom{\ell}{2} \frac{m}{\ell}$ edges are not covered by the S -matchings.

Hence the total number of edges that are not used in a full K_m can be bounded by

$$\left(\sigma + \eta \frac{m-1}{\ell-1} \right) \binom{\ell}{2} \frac{m}{\ell}.$$

It is left to establish a similar estimate for the other K_m 's in the K_m -factor. Recall that we declared the whole K_m -factor to be full as soon as one K_m was full. Since all components of any $\mathcal{C}_i \subseteq F_i$ have bounded maximum degree Δ , in each step up to at most $\frac{m\Delta}{2}$ edges are reserved in any K_m of the current K_m -factor. Owing to the balanced selection of the K_m 's within the current K_m -factor (see Step ii)) we have that the number of used edges over all K_m 's in $\mathcal{D}_{\text{current}}^{m,N}$ differs by at most $\frac{m\Delta}{2}$. Consequently, the number of unused edges in any K_m at the point when the K_m -factor is declared full is at most

$$\left(\sigma + \eta \frac{m-1}{\ell-1} \right) \binom{\ell}{2} \frac{m}{\ell} + \frac{m\Delta}{2}.$$

Using this estimate for all $\binom{N}{2} / \binom{m}{2}$ of the K_m in the K_m -decomposition of K_N leads to a total of unused edges of at most

$$\left(\sigma \frac{\ell-1}{m-1} + \eta + \frac{\Delta}{m-1} \right) \binom{N}{2} < 2\eta \binom{N}{2},$$

where we used (6) and $\Delta < \sigma$. Furthermore, since by $N \geq (1 + \frac{\xi}{2})n$ we have

$$\binom{n}{2} + 2\eta \binom{N}{2} \leq \binom{N}{2},$$

we have shown that we do not run out of K_m -factors and, hence, the assignment procedure yields a preliminary packing of $\{\mathcal{C}_i\}_{i \in [n]}$.

For the proof of Lemma 10 we have to show not only that there exists such a packing but also that there is a balanced one. This will be the focus of the next phase.

5.4.2. The balancing phase. In the assignment phase we have constructed a preliminary packing \mathbf{f} of $\{\mathcal{C}_i\}_{i \in [n]}$ into the K_m -factors of K_N as described in Section 5.1. We now construct a (ξn) -balanced packing \mathbf{h} by the following random process consisting of two parts. Firstly, we randomly permute the $\frac{N}{m}$ K_m 's in each K_m -factor independently, and we will denote the resulting packing by \mathbf{g} . Secondly, for each K_m , we pick a random permutation of its vertices. As we already noted in Section 5.3, any such permutation yields a packing of $\{\mathcal{C}_i\}_{i \in [n]}$ into K_N .

It is left to show that with positive probability each vertex v of K_N has boundary degree with respect to \mathbf{h} bounded by ξn . Recall from Definition 9 that the *boundary degree* with respect to \mathbf{f} of a vertex v is defined by

$$d_{\mathbf{f}}^{\partial}(v) = |\{i \in [n]: f_i^{-1}(v) \in \partial\mathcal{C}_i\}|.$$

For a K_m of the K_m -decomposition of K_N and a vertex v of K_m we consider the *relative boundary degree*

$$d_{\mathbf{f}}^{\partial}(v, K_m) = |\{i \in [n]: f_i \text{ assigns some components of } \mathcal{C}_i \text{ to } K_m \text{ and } f_i^{-1}(v) \in \partial\mathcal{C}_i\}|$$

Clearly, $\sum d_{\mathbf{f}}^{\partial}(v, K_m) = d_{\mathbf{f}}^{\partial}(v)$, where the sum runs over all K_m from the K_m -decomposition of K_N that contain v . For each K_m we define its *label* as the monotone sequence of the boundary degrees of its vertices. Since only one isomorphism type $S \in \mathcal{S}$ appears in each K_{ℓ} -factor, we can bound the number of possible labels by $\sigma^{(m-1)/(\ell-1)} < 2^m$ (see (3) and the choice of $\ell > s^2$). We call such a label *common* if at least $\frac{\eta}{2^m} \frac{N(N-1)}{m(m-1)}$ K_m 's have that label and *rare* otherwise. Note that the total number of K_m 's having a rare label is bounded by $\eta \frac{N(N-1)}{m(m-1)}$.

We use these labels to show that each vertex in K_N hosts roughly the same amount of boundary vertices. For that we first prove that an arbitrary vertex is incident to approximately the *expected* number of K_m 's of a given common label. For a vertex v of K_N and a common label A we denote by $X^{v,A}$ the number of K_m 's incident to v that have label A . Moreover, let $\alpha_k(A)$ be the number of K_m 's in the K_m -factor $\mathcal{D}_k^{m,N}$ that have label A and define

$$\alpha(A) = \sum_{k=1}^{\frac{N-1}{m-1}} \alpha_k(A).$$

Note that $X^{v,A}$ is the sum of $\frac{N-1}{m-1}$ indicator variables $X_k^{v,A}$, where $X_k^{v,A} = 1$ if the K_m from the K_m -factor $\mathcal{D}_k^{m,N}$ adjacent to v has label A . The probability that this happens is then given by $\frac{\alpha_k(A)}{N/m}$. By applying Chernoff's inequality ((2.9) in [9]) we obtain

$$\mathbb{P}(|X^{v,A} - \mathbb{E}X^{v,A}| > \eta \mathbb{E}X^{v,A}) < 2 \exp\left(-\frac{\eta^2 \mathbb{E}X^{v,A}}{3}\right) < 2 \exp\left(-\frac{\eta^2 m}{3 N} \alpha(A)\right).$$

Consequently, the probability that one of the common labels appears too many or too few times among the K_m 's incident to some vertex is bounded by

$$\sum_{v \in V(K_N)} \sum_{A \text{ common}} 2 \exp\left(-\frac{\eta^2 m}{3 N} \alpha(A)\right) < N 2^{m+1} \exp\left(-\frac{\eta^3 (N-1)}{2^m \cdot 3(m-1)}\right) < 1.$$

Therefore, with positive probability, all vertices are *balanced* in the sense that the occurrences of every common label among the K_m 's incident to each vertex roughly agree in proportion with the occurrences of that label in the decomposition.

We fix such permutation of the K_m 's and the corresponding numbers $X^{v,A}$ for every vertex v and every label A . Let \mathbf{g} be the corresponding packing of $\{\mathcal{C}_i\}_{i \in [n]}$ into K_N . As a consequence, we get that for every vertex v the number of K_m 's with common labels attached to it satisfies

$$\begin{aligned} \sum_{A \text{ common}} X^{v,A} &\geq \sum_{A \text{ common}} (1-\eta) \frac{m}{N} \alpha(A) = (1-\eta) \frac{m}{N} \sum_{A \text{ common}} \alpha(A) \\ &\geq (1-\eta) \frac{m}{N} (1-\eta) \frac{N(N-1)}{m(m-1)} = (1-\eta)^2 \frac{N-1}{m-1} \geq (1-2\eta) \frac{N-1}{m-1}. \end{aligned}$$

We also obtain an upper bound on the number of K_m 's with rare labels for every vertex v :

$$\sum_{A \text{ rare}} X^{v,A} = \frac{N-1}{m-1} - \sum_{A \text{ common}} X^{v,A} \leq 2\eta \frac{N-1}{m-1}. \quad (10)$$

Next we show that randomly permuting the vertices of each K_m in the K_m -decomposition of K_N for the random packing \mathbf{g} ensures that the boundary degrees in each K_m are evenly distributed. Let $d_{\mathbf{h}}^{\partial}(v, A)$ be the sum of the boundary degrees of the vertex v within the K_m 's labelled by A and containing v . Clearly,

$$d_{\mathbf{h}}^{\partial}(v) = \sum_A d_{\mathbf{h}}^{\partial}(v, A).$$

For a moment we ignore the K_m 's with rare labels, since owing to (10) their contribution will be negligible, and consider only those that have a common label. We denote by $A(i)$ the i -th element of the degree sequence A and set $\beta(A) = \frac{1}{m} \sum_{i=1}^m A(i)$ as the average degree in A .

We first show that for a vertex v of K_N and a common label A , $d_{\mathbf{h}}^{\partial}(v, A)$ is in the range $(1 \pm \eta)\beta(A)X^{v,A}$ with high probability. Let $Y_i^{v,A}$ be the number of K_m 's labelled by A in which v gets boundary degree $A(i)$. By applying Chernoff's inequality we obtain

$$\mathbb{P} \left(\left| Y_i^{v,A} - \frac{X^{v,A}}{m} \right| > \eta \frac{X^{v,A}}{m} \right) < 2 \exp \left(-\frac{\eta^2}{3} \frac{X^{v,A}}{m} \right)$$

for every $i \in [m]$. This implies that with probability $1 - 2m \exp \left(-\frac{\eta^2}{3} \frac{X^{v,A}}{m} \right)$ we have

$$d_{\mathbf{h}}^{\partial}(v, A) = \sum_{i=1}^m A(i) Y_i^{v,A} = \sum_{i=1}^m A(i) (1 \pm \eta) \frac{X^{v,A}}{m} = (1 \pm \eta) \beta(A) X^{v,A}.$$

By summing over all common labels, we have that with positive probability there exist permutations for every K_m of the K_m -decomposition of K_N for which all vertices have

roughly the expected boundary degree. More precisely, the probability that there exists a misbehaving vertex is bounded by

$$\sum_{v \in V(K_N)} \sum_{A \text{ common}} 2m \exp\left(-\frac{\eta^2}{3} \frac{X^{v,A}}{m}\right) < N 2^{m+1} m \exp\left(-\frac{\eta^3}{2^m \cdot 3} (1-\eta) \frac{N-1}{m(m-1)}\right) < 1.$$

Therefore, the contribution of the K_m 's with common labels for each vertex v is at most

$$\begin{aligned} \sum_{A \text{ common}} d_{\mathbf{h}}^\partial(v, A) &\leq \sum_{A \text{ common}} (1+\eta)\beta(A)X^{v,A} \leq \sum_{A \text{ common}} (1+\eta) \left[\beta(A)(1+\eta) \frac{m}{N} \alpha(A) \right] \\ &= (1+\eta)^2 \frac{1}{N} \sum_{A \text{ common}} (m\beta(A)\alpha(A)) \leq (1+\eta)^2 \frac{1}{N} \sum_{i=1}^n |\partial \mathcal{C}_i| \\ &\stackrel{(1)}{\leq} (1+\eta)^2 \frac{1}{N} \xi n^2. \end{aligned}$$

Owing to (10), a vertex can be incident to at most $2\eta \frac{N-1}{m-1}$ K_m 's with rare labels. Since no component of any \mathcal{C}_i consists of a single isolated vertex (see ii. in Definition 5), the largest relative boundary degree of any vertex in such a K_m can be at most $m-1$ and we infer that

$$d_{\mathbf{h}}^\partial(v) \leq (1+\eta)^2 \frac{1}{N} \xi n^2 + 2\eta \frac{N-1}{m-1} (m-1) < \left(\frac{(1+\eta)^2 \xi}{1+\xi/2} + 2\eta(1+\xi) \right) n < \xi n \quad (11)$$

for every $v \in V(K_N)$, thus proving Lemma 10. \square

6. PACKING THE SEPARATORS

In this section we prove Lemma 11. The Lemma asserts that a balanced packing of $\{\mathcal{C}_i\}_{i \in [n]}$ into $K_{(1+\xi)n}$ can be extended to a packing of $\{F_i\}_{i \in [n]}$ in $K_{(1+\varepsilon)}$. For that we have to show that we can embed the separators $\{U_i\}_{i \in [n]}$ in an appropriate way. Roughly speaking, we will show that a simple greedy strategy will work in here.

6.1. Proof of Lemma 11. Given ε and Δ , set

$$\xi = \frac{\varepsilon}{12\Delta^2} \quad \text{and} \quad \delta = \frac{\varepsilon^2}{72\Delta^2}.$$

Let $s \in \mathbb{N}$ and let \mathcal{G} be a (δ, s) -separable family. For some sufficiently large n let $\mathcal{F} = (F_1, \dots, F_n)$ be a (\mathcal{G}, n, Δ) -sequence and suppose that there exists a (ξn) -balanced packing of the component graphs $\{\mathcal{C}_i\}_{i \in [n]}$ into a clique of order $(1+\xi)n$. Fix a partition $X \dot{\cup} Y$ of the vertex set of $K_{(1+\varepsilon)n}$, where $|X| = (1+\xi)n$, and denote by K_X , K_Y , and $K_{X,Y}$ the complete subgraphs induced on X and on Y , and the complete bipartite subgraph between X and Y , respectively. Let $\mathbf{h} = \{h_i\}_{i \in [n]}$ with

$$h_i: V(F_i) \rightarrow X$$

be a (ξn) -balanced packing of $\{\mathcal{C}_i\}_{i \in [n]}$ into K_X . We shall use K_Y to embed $\{U_i\}_{i \in [n]}$, and $K_{X,Y}$ for the necessary connections. It is easy to see that if the following conditions are satisfied then the resulting map is a packing of \mathcal{F} into $K_{(1+\varepsilon)n}$:

- (1) for every $i \in [n]$, the vertices of U_i are mapped injectively into Y ;
- (2) each edge in $K_{X,Y}$ is used at most once;
- (3) each edge in K_Y is used at most once.

Note that we will embed $\sum_{i \in [n]} |U_i| \leq \delta n^2$ vertices into Y , therefore some vertices in Y will be used at least $\frac{\sum_{i \in [n]} |U_i|}{|Y|} \leq \frac{\delta n^2}{|Y|}$ times. However, we will ensure that each vertex in Y is used at most $3 \frac{\delta n^2}{|Y|}$ times. The packing of \mathcal{F} into $K_{(1+\varepsilon)n}$ will be expressed by a family of functions $\bar{\mathbf{h}} = \{\bar{h}_i\}_{i \in [n]}$ with

$$\bar{h}_i: V(F_i) \rightarrow X \dot{\cup} Y$$

where \bar{h}_i extends h_i from $V(\mathcal{C}_i)$ to $V(F_i)$. For a vertex $v \in V(\mathcal{C}_i)$, we set $\bar{h}_i(v) = h_i(v) \in X$ for any $i \in [n]$. For the vertices in the separators $\{U_i\}_{i \in [n]}$ we will fix their image $\bar{h}_i(v)$ in Y one by one in a greedy way, starting with vertices of U_1 .

At each step we embed a vertex $u \in U_i$ into Y , assuming that all vertices of U_j with $j < i$ and possibly some (at most $|U_i| - 1 < \delta n$) vertices of U_i were already embedded. Let $N_{\mathcal{C}_i}(u)$ be the neighbourhood of u in \mathcal{C}_i , and $N_{U_i}(u)$ the neighbourhood of u in U_i both of size at most Δ . Suppose so far we made sure that every vertex in Y was used at most $3 \frac{\delta n^2}{|Y|}$ times. We will embed u in such a way that (1), (2), and (3) are obeyed (see (1*), (2*), and (3*) below), and afterwards each vertex of Y is still used at most $3 \frac{\delta n^2}{|Y|}$ times. This will show that \mathbf{h} can be extended to a packing $\bar{\mathbf{h}}$ of \mathcal{F} and conclude the proof. Having this in mind we note:

- (1*) The vertices of U_i have to be embedded injectively into Y and, hence, up to at most $|U_i| - 1 < \delta n$ vertices of Y may not be used for the embedding of u .
- (2*) Since every edge in $K_{X,Y}$ can be used at most once, we require $\bar{h}_i(u) \neq \bar{h}_j(u')$ for every vertex $u' \in U_j$ with $\bar{h}_j(N_{\mathcal{C}_j}(u')) \cap \bar{h}_i(N_{\mathcal{C}_i}(u)) \neq \emptyset$. Let $x \in \bar{h}_i(N_{\mathcal{C}_i}(u))$. Owing to the (ξn) -balancedness of the packing $\{h_i\}_{i \in [n]}$, x hosts at most ξn vertices from $\bigcup_{k \in [n]} \partial \mathcal{C}_k$ and each of them has at most Δ neighbours in some U_k for $k \in [n]$. Assuming that all of them have already been embedded into Y , we obtain at most $\Delta \xi n$ forbidden vertices for each of the up to at most Δ neighbours of u in \mathcal{C}_i . Hence, the total number of forbidden options for $\bar{h}_i(u)$ in Y is at most $\Delta^2 \xi n$.
- (3*) Note that K_Y also hosts the edges contained in the separator U_i and every edge of K_Y may be used at most once. Suppose that there exists a vertex u' from U_j with $j < i$ such that $\bar{h}_i(N_{U_i}(u)) \cap \bar{h}_j(N_{U_j}(u')) \neq \emptyset$. Then $\bar{h}_i(u)$ must avoid $\bar{h}_j(u')$ for any such u' , because at least one edge between this vertex and the image of the neighbours of u is already used. Since by our assumption every vertex in the set $\bar{h}_i(N_{U_i}(u))$ hosts at most $3 \frac{\delta n^2}{|Y|}$

vertices embedded so far, and since $\Delta(F_j) \leq \Delta$, there are at most $\Delta \cdot 3 \frac{\delta n^2}{|Y|} |N_{U_i}(u)| \leq 3\Delta^2 \delta n^2 / |Y|$ such restrictions.

Since up to now every vertex $y \in Y$ was used at most $3 \frac{\delta n^2}{|Y|}$ times for the embedding, by denoting with $Y_u \subseteq Y$ the set of candidates for the embedding of u , we obtain

$$|Y_u| \geq |Y| - (\delta n + \Delta^2 \xi n + 3\Delta^2 \frac{\delta n^2}{|Y|}) \geq |Y| - \frac{\varepsilon}{4} n > \frac{|Y|}{2}.$$

Since we have to embed at most $\sum_{i \in [n]} |U_i| \leq \delta n^2$ vertices in total, at any time some vertex $y \in Y_u$ was used at most

$$\frac{\delta n^2}{|Y|/2} < 3 \frac{\delta n^2}{|Y|} - 1$$

times, and this vertex we choose for $\bar{h}_i(u)$. We have thus shown that at each round we can always pick one vertex in Y such that all the edges needed to connect the vertex we want to embed to all its neighbour are available and it was used before at most $3 \frac{\delta n^2}{|Y|} - 1$ times. This completes the proof of the lemma. \square

REFERENCES

1. N. Alon, P. Seymour, and R. Thomas, *A separator theorem for nonplanar graphs*, J. Amer. Math. Soc. **3** (1990), no. 4, 801–808.
2. N. Alon and J. H. Spencer, *The probabilistic method*, third ed., Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., Hoboken, NJ, 2008, With an appendix on the life and work of Paul Erdős.
3. N. Alon and R. Yuster, *Every H -decomposition of K_n has a nearly resolvable alternative*, European J. Combin. **21** (2000), no. 7, 839–845.
4. J. Böttcher, J. Hladký, D. Piguet, and A. Taraz, *An approximate version of the tree packing conjecture*, Israel Journal of Mathematics, to appear.
5. P. Frankl and V. Rödl, *Near perfect coverings in graphs and hypergraphs*, European J. Combin. **6** (1985), no. 4, 317–326.
6. R. L. Graham, M. Grötschel, and L. Lovász (eds.), *Handbook of combinatorics. Vol. 1, 2*, Elsevier Science B.V., Amsterdam; MIT Press, Cambridge, MA, 1995.
7. A. Gyárfás and J. Lehel, *Packing trees of different order into K_n* , Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely, 1976), Colloq. Math. Soc. János Bolyai, vol. 18, North-Holland, Amsterdam, 1978, pp. 463–469.
8. A. M. Hobbs, *Packing trees*, Proceedings of the Twelfth Southeastern Conference on Combinatorics, Graph Theory and Computing, Vol. II (Baton Rouge, La., 1981), vol. 33, 1981, pp. 63–73.
9. S. Janson, T. Łuczak, and A. Ruciński, *Random graphs*, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley-Interscience, New York, 2000.
10. N. Pippenger and J. Spencer, *Asymptotic behavior of the chromatic index for hypergraphs*, J. Combin. Theory Ser. A **51** (1989), no. 1, 24–42.

FACHBEREICH MATHEMATIK, UNIVERSITÄT HAMBURG, BUNDESSTRASSE 55, D-20146 HAMBURG, GERMANY

E-mail address: {`silvia.messuti|schacht`}@math.uni-hamburg.de

DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE, EMORY UNIVERSITY, ATLANTA, GA 30322, USA

E-mail address: `rod1@mathcs.emory.edu`