PACKING MINOR CLOSED FAMILIES OF 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 $\sum_{i=1}^n e(T_i) \le \binom{n}{2}$ and with bounded maximum degree, can be packed into the complete graph on $(1+o(1))n$ vertices. We generalize 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.

A packing of a sequence of graphs $\mathcal{F} = (F_1, \dots, F_n)$ into a graph H is a collection of edge-disjoint subgraphs $H_1, \dots, H_n \subseteq H$, such that H_i is isomorphic to F_i for every $i \in [n]$. A well-known conjecture of Gyárfás [4] states that a sequence of trees $\mathcal{T} = (T_1, \dots, T_n)$, where $v(T_i) = i$ for every $i \in [n]$, packs into K_n. Note that the sum of the edges over $\mathcal{T}$ is precisely $\binom{n}{2}$, hence the packing of $\mathcal{T}$ should use all the edges of K_n. A restricted approximate version where the host graph is a clique on $(1+o(1))n$ vertices and the trees have at most n vertices, bounded maximum degree, and the sum over all edges is at most $\binom{n}{2}$, was proved by Böttcher, Hladký, Piguet, and Taraz [3]. We extend this result to sequences of graphs from any non-trivial minor closed family.

Theorem 1. *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 $n \ge n_0$ the following holds. If $\mathcal{F} = (F_1, \dots, F_n)$ is a sequence of graphs from $\mathcal{G}$, each having order at most n and maximum degree at most Δ, such that $\sum_{i=1}^n e(F_i) \le \binom{n}{2}$, then $\mathcal{F}$ packs into $K_{(1+\varepsilon)n}$.*

The main idea in the proof is to remove a small separator from each graph, in such a way that all components have small constant size, then pack the components into a large clique contained into $K_{(1+\varepsilon)n}$ and use the remaining vertices for the separators. In fact we prove a more general result and derive Theorem 1 from that, and the Separator Theorem of Alon, Seymour, and Thomas [1]. We define a (δ, s)-separable family as a set of graphs each having the property that by removing δ proportion of the vertices, the resulting components have size at most s.

Theorem 2. *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 $n \ge n_0$ the following holds. If $\mathcal{F} = (F_1, \dots, F_n)$ is a sequence of graphs from $\mathcal{G}$ each having order at most n and maximum degree at most Δ, such that $\sum_{i=1}^n e(F_i) \le \binom{n}{2}$, then $\mathcal{F}$ packs into $K_{(1+\varepsilon)n}$.*

As mentioned above, Theorem 1 easily follows from Theorem 2. In fact, the Separator Theorem states that for any minor closed family $\mathcal{G}$ there exists a constant c such that a graph $G \in \mathcal{G}$ of order n has a separator of order at most $cn^{1/2}$ and the components have size at most $n/2$. Applying such result recursively to each component, for $i > 0$ iterations, leads to a separator $U \subseteq V(G)$ such that

$$|U| \le cn^{1/2} + 2c\left(\frac{n}{2}\right)^{1/2} + \cdots + 2^{i-1}c\left(\frac{n}{2^{i-1}}\right)^{1/2} = cn^{1/2} \cdot \frac{\sqrt{2}^{\,i} - 1}{\sqrt{2} - 1} < 3cn^{1/2}\, 2^{i/2}$$

and each component of $G - U$ has order at most $n/2^i$. For given $\delta > 0$ we can apply this with $i = \left\lfloor 2\log_2\left(\frac{\delta n^{1/2}}{3c}\right)\right\rfloor$ and obtain a separator U of order at most δn, and a set of components all of

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).

which have order at most $18c^2/\delta^2$. This shows that the minor-closed family $\mathcal{G}$ is $\left(\delta, \frac{18c^2}{\delta^2}\right)$-separable for any $\delta > 0$. Hence, we can apply Theorem 2, and Theorem 1 follows.

In the rest of the abstract we give a sketch of the proof of Theorem 2 in which we ignore the choice of the parameters and focus on the construction of the packing. So we assume that for ε and Δ we are given a sequence $\mathcal{F}$ (from some appropriate (δ, s)-separable family $\mathcal{G}$) that satisfies the assumptions of the theorem, and that n is sufficiently large. We consider each graph F_i of the sequence together with a fixed (δ, s)-separation $(U_i, \mathcal{C}_i)$, where $U_i \subseteq V(F_i)$ is the separator and has size $|U_i| \leq \delta v(F_i)$, and $\mathcal{C}_i$ is the set of the components of $F_i \backslash U_i$, where each component C has size at most s. We define the boundary ∂C of a component C as the subset of the vertices that are adjacent to the separator, and more generally the boundary of a set of components as the union of the boundaries of all its components, and the boundary of $\mathcal{F}$ as the union of the boundaries of all its component sets.

The vertices of the host graph $K_{(1+\varepsilon)n}$ will be split into a large part X of order $(1 + o(1))n$ and a small part $Y = V \backslash X$ of order $o(n)$. We will pack the graphs $\{\mathcal{C}_i\}_{i\in[n]}$ into the clique K_X spanned on X and the sets $\{U_i\}_{i\in[n]}$ into the clique K_Y spanned on Y. For this, we shall ensure that all vertices from the boundary $\partial\mathcal{C}_i$ will be appropriately connected to the vertices of the separator U_i. Since every edge of the complete bipartite graph $K_{X,Y}$ induced by X and Y can be used only once in the packing, we will make sure that the packing of $\{\mathcal{C}_i\}_{i\in[n]}$ will be balanced, in the sense that each vertex of X will only be the target for not too many boundary vertices.

Packing the components. The packing of $\{\mathcal{C}_i\}_{i\in[n]}$ consists of two steps: the assignment and the embedding. In the assignment phase we group the components by isomorphism types, decompose K_X into such graphs, and assign each component to one of its copies in the decomposition. In the embedding phase we actually "place" the copies of the components somewhere into K_X, i.e., we assign the copies of the components to the vertices of K_X. We will do this randomly, and show that with positive probability the resulting packing is balanced.

As we mentioned, the idea is to decompose K_X into smaller graphs isomorphic to the components in $\{\mathcal{C}_i\}_{i\in[n]}$ and assign each component to one its copies in the decomposition. However, two components from the same graph have to be vertex disjoint by definition, so we need a decomposition that can be partitioned into (possibly perfect) matchings. For that we apply a result by Ray-Chaudhuri and Wilson [5] that states that for $m \geq 2$ and sufficiently large n, if the necessary divisibility conditions are met, then there exists a resolvable K_m-decomposition of K_n, i.e., a decomposition that can be partitioned into K_m-factors. For general G, resolvable decompositions do not necessarily exist. However, due to Alon and Yuster [2], we can decompose a K_n into G in such a way that almost all matchings contain a large number of copies of G, if the divisibility conditions are satisfied and n is sufficiently large.

Let $\mathcal{S} = \{S_1, \ldots, S_\sigma\}$ be the set of graphs on at most s vertices and maximum degree at most Δ isomorphic to some component in $\{\mathcal{C}_i\}_{i\in[n]}$, with $e(S_1) \geq \cdots \geq e(S_\sigma)$. We assume that X has size $(1 + \xi)n$ for some appropriately chosen ξ. We will decompose a subgraph K_N of K_X where $(1 + \frac{\xi}{2})n \leq N \leq (1 + \xi)n$ and there exist m and ℓ with $s \ll \ell \ll m \ll N$ for which K_N has a resolvable K_m-decomposition $\mathcal{D}^{m,N}$, K_m has a resolvable K_ℓ-decomposition $\mathcal{D}^{\ell,m}$, and K_ℓ can be decomposed into S in such a way that at least $(1-\zeta)\frac{(\ell-1)v(S)}{2e(S)}$ matchings of the decomposition $\mathcal{D}^{S,\ell}$ contain at least $(1-\zeta)\frac{\ell}{v(S)}$ copies of S, where S is any graph from $\mathcal{S}$ and ζ is an appropriately chosen constant. Note that using the same resolvable decomposition $\mathcal{D}^{\ell,m}$ in all copies of K_m in $\mathcal{D}^{m,N}$ yields a resolvable K_ℓ-decomposition of K_N which we denote by $\mathcal{D}^{\ell,N}$. The copies of K_ℓ in such decomposition may be partitioned into K_ℓ-factors in many ways. In fact, given a K_m-factor of K_N, we may select a "different" K_ℓ-factor inside each K_m independently and obtain a K_ℓ-factor of K_N this way.

We shall pack the component graphs $\{\mathcal{C}_i\}_{i\in[n]}$ into K_N one by one in such a way that each $\mathcal{C}_i$ is completely contained in one of the K_ℓ-factors of K_N. Roughly speaking, for each $\mathcal{C}_i$ we group the components by isomorphism types, and for every such isomorphism type $S \in \mathcal{S}$ we allocate an S-matching in sufficiently many copies of K_m. Since for different graphs in $\{\mathcal{C}_i\}_{i\in[n]}$ we may need to use the same K_m for hosting different graphs from $\mathcal{S}$, we decompose K_m into copies of K_ℓ. This way, given different S and S' from $\mathcal{S}$, we may allocate an S-matching in K_m by using an S-decomposition into one K_ℓ-factor, and an S'-matching in the same K_m by using an S'-decomposition in a different K_ℓ-factor. We proceed this way until essentially all edges of the given K_m-factor are used. At this point we switch to the next K_m-factor in K_N and proceed the same way. In order to keep the amount of unused edges small during such a change of K_m-factors, we have to make sure that all K_m's in such a factor are "populated" by component graphs $\mathcal{C}_i$ in a balanced way. This will be achieved by choosing the copies of K_m for each $\mathcal{C}_i$ among those that have the least number of used edges. We now describe the assignment of the components of a graph in detail.

For the set of components $\mathcal{C}_i$ of F_i and for a graph $S \in \mathcal{S}$, we define $M_i(S)$ as the number of copies of K_m we shall reserve for embedding all the components isomorphic to S from $\mathcal{C}_i$ into K_N. Note that we ignore the S-matchings that contain less than $(1-\zeta)\frac{\ell}{v(S)}$ copies of S. At any time we work in a fixed K_m-factor in which the K_m's are considered to be ordered starting from the one in which the least number of edges has been allocated to some of the previous component graphs. Moreover, in each K_m we have for every $S \in \mathcal{S}$ a "current" K_ℓ-factor in which all K_ℓ's are S-decomposed. We assign the components of $\mathcal{C}_i$ to $M_i = \sum_{s=1}^{\sigma} M_i(S_s)$ of these K_ℓ-factors each from a different K_m as described below.

We start by picking the first M_i copies of K_m in the current K_m-factor and assign each of them a graph from $\mathcal{S}$ isomorphic to some component of $\mathcal{C}_i$. Recall that the K_m's are ordered from the one with the most number of available edges to that with the most number of used edges, and the graphs in $\mathcal{S}$ are ordered from the densest to the sparsest. Hence, by assigning S_1 to the first $M_i(S_1)$ copies of K_m, S_2 to the next $M_i(S_2)$ copies, and so on, we maintain an "approximately balanced" assignment to the current K_m-factor. Once we decide which S goes to which K_m, we look at the current S-decomposed K_ℓ-factor in the respective K_m's and allocate a large S-matching to which we assign the components from $\mathcal{C}_i$ that are isomorphic to S.

Before we repeat the procedure for $\mathcal{C}_{i+1}$, we update our workspace as follows. For every K_ℓ-factor we just used for the components from $\mathcal{C}_i$ isomorphic to some S we check whether it contains another large S-matching that has not been allocated yet. Note that each time we consider a K_ℓ-factor for the assignment we use all the K_ℓ's contained in it, hence the S-matchings get exhausted in all the K_ℓ's at the same moment. If all the large S-matchings have been used, then we pick a new K_ℓ-factor, if available, in the same K_m. To each newly selected K_ℓ we apply $\mathcal{D}^{S,\ell}$, and this becomes the current K_ℓ-factor for S in this K_m. If we are in the situation that there is no available K_ℓ-factor in at least one of the K_m's, then we consider the K_m's of this K_m-factor in K_N as completely used, and we will show that the "balancing procedure" described above ensures that most K_m's are "almost full". At this point we replace our current K_m-factor with a new one, and in each K_m we set up current K_ℓ-factors for each $S \in \mathcal{S}$.

The procedure above produces an assignment of the components in $\{\mathcal{C}_i\}_{i\in[n]}$ to their copies inside K_N such that all components from each $\mathcal{C}_i$ are vertex disjoint. It is left to show that there is enough room for all the component graphs, i.e., that the edges we waste in the assignment procedure plus those contained into the S-matchings that were not used because they contained too few copies of S are less than the additional edges we allow by considering K_N instead of K_n.

In order to determine the number of edges wasted in the assignment procedure, observe that a K_m is considered to be full when for some S there is no large S-matching available in the current K_ℓ-factor and all other K_ℓ-factors in that K_m have been used, or are currently reserved for some other $S' \in \mathcal{S}$. At this moment, in the worst case all current K_ℓ-factors are empty, hence the

waste is at most $\sigma\frac{m}{\ell}\binom{\ell}{2}$. Since we pick the least used K_m's, we may assume that in every other K_m of the current K_m-factor we have the same waste, yielding a total waste of $\frac{N}{m}\sigma\frac{m}{\ell}\binom{\ell}{2}$ pairs in the K_m-factor, and hence $\sigma\frac{\ell-1}{m-1}\binom{N}{2}$ edges in the whole K_N. As for the edges contained in small S-matchings, one can show that their total number is at most $(\zeta-\zeta^2)\binom{N}{2}$. By our choice of m and ζ we have that the total number of unused edges in K_N

$$\sigma\frac{\ell-1}{m-1}\binom{N}{2}+(\zeta-\zeta^2)\binom{N}{2}$$

is smaller than $\binom{N}{2}-\binom{n}{2}$. This shows that all component graphs in $\mathcal{F}$ fit into K_N.

For the actual embedding, i.e., the assignment of the copies of the components to the vertices of K_N, we have to make sure that each vertex of K_N hosts at most ξn vertices from the boundary of $\mathcal{F}$. For that we will first pick a random permutation of the K_m's in each K_m-factor, and then a random permutation of the vertices inside each K_m.

We define a labelling of the K_m's in the assignment by assigning each K_m the (unordered) sequence of the boundary degrees of its vertices, i.e., for each vertex in the K_m we count how many times it was used for the embedding of a boundary vertex. For some appropriately chosen constant η we say that a label is *common* if at least $\frac{\eta}{2^m}\frac{N(N-1)}{m(m-1)}$ K_m's have that label and *rare* otherwise. We shall also use η as the multiplicative error in the following applications of Chernoff's inequality.

Given the K_m-decomposition of K_N, we look at one vertex v and choose the K_m's incident to v by randomly picking one K_m in each K_m-factor. Some applications of Chernoff's inequality show that with positive probability, for all vertices we have 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. Since all but at most an η-proportion of the K_m's in the decomposition have a common label, we obtain that the number of K_m's with common labels attached to each vertex v is at least $(1-\eta)^2\frac{N-1}{m-1}$, and the number of K_m's with rare labels for v is at most $2\eta\frac{N-1}{m-1}$.

Next we shall also permute the vertices inside each K_m so that the boundary degrees in each K_m will be equally distributed. For a vertex v and a common label A we pick a random value in the boundary degree sequence of each K_m labelled with A and by applying Chernoff's inequality we can show that the sum of boundary degrees of v in all these K_m's is concentrated around its expectation. By summing over all common labels, we have that with positive probability there exists a permutation for which all vertices have boundary degree at most $(1+\eta)^2\frac{1}{N}\frac{\xi}{2}n^2$. By adding the largest possible boundary degree, i.e., m for at most $2\eta\frac{N-1}{m-1}$ K_m's with rare labels, we obtain that each vertex of K_N hosts at most

$$(1+\eta)^2\frac{1}{N}\frac{\xi}{2}n^2+2\eta\frac{N-1}{m-1}m<\left(\frac{(1+\eta)^2}{1+\xi/2}\frac{\xi}{2}+2\eta(1+\xi)\right)n<\xi n$$

vertices from the boundary of $\mathcal{F}$.

Packing the separators. Having constructed a packing of the components where each vertex of X hosts at most ξn vertices from the boundary of $\mathcal{F}$, it is left to show that we can use just a few more vertices to embed the separators. In order to obtain a packing of $\mathcal{F}$ we have to satisfy the following conditions:

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

We embed the separators one by one, vertex by vertex. In the following we describe the embedding of an arbitrary vertex u from an arbitrary separator U_i, assuming that all vertices of U_j with $j<i$ and up to at most $|U_i|\leq\delta n$ vertices of U_i were already embedded, and that 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 the constraints discussed above will be obeyed, and afterwards each vertex of Y is still used at most $3\frac{\delta n^2}{|Y|}$ times. For that we collect the restrictions given by the previous embeddings.

(1) Since the vertices of U_i have to be embedded injectively into Y, up to at most $|U_i| \leq \delta n$ vertices in Y might be forbidden for the embedding of u.

(2) Every edge in $K_{X,Y}$ can be used at most once, hence we have to embed u into a vertex y such that all edges between y and the vertices X' of X that embed the neighbours of u in $\mathcal{C}_i$ are available. Since u has at most Δ neighbours in $\mathcal{C}_i$, and each vertex in X' hosts at most ξn boundary vertices, each having at most Δ neighbours in $\{U_i\}_{i\in[n]}$, at most $\Delta^2\xi n$ more vertices might be forbidden.

(3) Also the edges in K_Y can be used at most once. This means that the target vertex for u must be chosen in such a way that all the edges to the vertices Y' that embed its neighbourhood in U_i have not been used yet. Note that each vertex in Y' may already host up to $3\frac{\delta n^2}{|Y|}$ vertices from $\{U_i\}_{i\in[n]}$, hence the embeddings of the neighbours of such vertices are also forbidden for the choice of u. Since u has at most Δ neighbours in U_i, and each vertex embedded into Y' has at most Δ neighbours in Y, this results in excluding at most $\Delta^2 3\frac{\delta n^2}{|Y|}$ additional vertices from Y.

An appropriate choice of the constants ensures that the set of candidates for the embedding of u is larger than $|Y|/2$. Since we have to embed at most $\sum_{i\in[n]} |U_i| \leq \delta n^2$ vertices in total, some vertex $y \in Y$ was used at most

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

times, and this vertex we choose for the embedding of u. We have thus shown that for each vertex of $\{U_i\}_{i\in[n]}$ we can pick a vertex in Y such that all the edges needed for the necessary connections are available, and hence obtain a packing of $\mathcal{F}$.

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 R. Yuster, *Every H-decomposition of K_n has a nearly resolvable alternative*, European J. Combin. **21** (2000), no. 7, 839–845.
3. J. Böttcher, J. Hladký, D. Piguet, and A. Taraz, *An approximate version of the tree packing conjecture*, manuscript.
4. 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.
5. D. K. Ray-Chaudhuri and R. M. Wilson, *The existence of resolvable block designs*, Survey of combinatorial theory (Proc. Internat. Sympos., Colorado State Univ., Fort Collins, Colo., 1971), North-Holland, Amsterdam, 1973, pp. 361–375.

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: rodl@mathcs.emory.edu