

Large-Treewidth Graph Decompositions and Applications

Chandra Chekuri*

Julia Chuzhoy †

ABSTRACT

Treewidth is a graph parameter that plays a fundamental role in several structural and algorithmic results. We study the problem of decomposing a given graph G into node-disjoint subgraphs, where each subgraph has sufficiently large treewidth. We prove two theorems on the tradeoff between the number of the desired subgraphs h , and the desired lower bound r on the treewidth of each subgraph. The theorems assert that, given a graph G with treewidth k , a decomposition with parameters h, r is feasible whenever $hr^2 \leq k/\text{poly log}(k)$, or $h^3r \leq k/\text{poly log}(k)$ holds. We then show a framework for using these theorems to bypass the well-known Grid-Minor Theorem of Robertson and Seymour in some applications. In particular, this leads to substantially improved parameters in some Erdos-Pósa-type results, and faster running times for algorithms for some fixed-parameter tractable problems.

1. INTRODUCTION

Let $G = (V, E)$ be an undirected graph. We assume that the reader is familiar with the notion of treewidth of a graph G , denoted by $\text{tw}(G)$. The main question considered in this paper is the following. Given an undirected graph G , and integer parameters $h, r < \text{tw}(G)$, can G be partitioned into h node-disjoint subgraphs G_1, \dots, G_h such that for each i , $\text{tw}(G_i) \geq r$? It is easy to see that for this to be possible, $hr \leq \text{tw}(G)$ must hold. Moreover, it is not hard to show examples of graphs G (such as constant-degree expanders), where even for $r = 2$, the largest number of node-disjoint subgraphs of G with treewidth at least $r = 2$ is bounded

*Dept. of Computer Science, University of Illinois, Urbana, IL 61801. chekuri@illinois.edu. Supported in part by NSF grant CCF-1016684 and by TTI Chicago on a sabbatical visit.

†Toyota Technological Institute, Chicago, IL 60637. Email: cjulia@ttic.edu. Supported in part by NSF CAREER grant CCF-0844872 and Sloan Research Fellowship.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

by $h = O\left(\frac{\text{tw}(G)}{\log(\text{tw}(G))}\right)$.¹ In this paper we prove the following two theorems, that provide sufficient conditions for the existence of a decomposition with parameters h, r .

Theorem 1.1 *Let G be any graph with $\text{tw}(G) = k$, and let h, r be any integers with $hr^2 \leq k/\text{poly log } k$. Then there is an efficient² algorithm to partition G into h node-disjoint subgraphs G_1, \dots, G_h such that $\text{tw}(G_i) \geq r$ for each i .*

Theorem 1.2 *Let G be any graph with $\text{tw}(G) = k$, and let h, r be any integers with $h^3r \leq k/\text{poly log } k$. Then there is an efficient algorithm to partition G into h node-disjoint subgraphs G_1, \dots, G_h such that $\text{tw}(G_i) \geq r$ for each i .*

We observe that the two theorems give different tradeoffs, depending on whether r is small or large. It is particularly useful in applications that the dependence is linear in one of the parameters. We make the following conjecture, that would strengthen and unify the preceding theorems.

Conjecture 1 *Let G be any graph with $\text{tw}(G) = k$, and let h, r be any integers with $hr \leq k/\text{poly log } k$. Then G can be partitioned into h node-disjoint subgraphs G_1, \dots, G_h such that $\text{tw}(G_i) \geq r$ for each i .*

Motivation and applications. The starting point for this work is the observation that a special case of Theorem 1.2, with $h = \Omega(\log^2 k)$, is a critical ingredient in recent work on poly-logarithmic approximation algorithms for routing in undirected graphs with constant congestion [8, 9, 5]. In particular, [8] developed such a decomposition for edge-disjoint routing, and subsequently [5] extended it to the node-disjoint case. However, in this paper, we are motivated by a different set of applications, for which Theorem 1.1 is more suitable. These applications rely on the seminal work of Robertson and Seymour [32], who showed that there is a large grid minor in any graph with sufficiently large treewidth. The theorem below, due to Robertson, Seymour and Thomas [31], gives an improved quantitative bound relating the size of the grid minor and the treewidth.

¹Consider a constant degree n -node expander G with girth $\Omega(\log n)$; the existence of such graphs can be shown by the probabilistic method. Let G_1, \dots, G_h be any collection of node-disjoint subgraphs of G of treewidth at least 2 each. Then each graph G_i must contain a cycle, and by the lower bound on the girth of G , $|V(G_i)| = \Omega(\log n)$, implying that $h = O(n/\log n)$. On the other hand $\text{tw}(G) = \Omega(n)$.

²In this paper we use the term efficient algorithm to refer to an algorithm that runs in randomized polynomial-time.

Theorem 1.3 (Graph-Minor Theorem [31]) *Let G be any graph, and g any integer, such that $\text{tw}(G) \geq 20^{2g^5}$. Then G contains a $g \times g$ grid as a minor. Moreover, if G is planar, then $\text{tw}(G) \geq 6g - 4$ suffices.*

Kawarabayashi and Kobayashi [21] obtained an improved bound of $2^{O(g^2 \log g)}$ on the treewidth required to ensure a $g \times g$ grid minor, and a further improvement to a bound of $2^{O(g \log g)}$ was recently claimed by Seymour [33].

Notice that Theorem 1.3 guarantees a grid minor of size sub-logarithmic in the treewidth k in general graphs, and of size $\Omega(k)$ in planar graphs. Demaine and Hajiaghayi [11] extended the linear relationship between the grid minor size and the treewidth to graphs that exclude a fixed graph H as a minor (the constant depends on the size of H , see [21] for an explicit dependence). A $g \times g$ grid has treewidth g , and it can be partitioned into h node-disjoint grids of size $r \times r$ each, as long as $r\sqrt{h} = O(g)$. Thus, in a general graph G of treewidth k , the Grid-Minor Theorem currently only guarantees that for any integers h, r with $hr^2 = O(\log^{2/5} k)$, there is a partition of G into h node-disjoint subgraphs of treewidth at least r each. Robertson et al. [31] observed that, in order for G to contain a $g \times g$ grid as a minor, its treewidth may need to be as large as $\Omega(g^2 \log g)$, and they suggest that this may be sufficient. Demaine et al. [12] conjecture that the treewidth of $\Theta(g^3)$ is both necessary and sufficient.

The existence of a polynomial relationship between the grid-minor size and the graph treewidth is a fundamental open question, that appears to be technically very challenging to resolve. Our work is motivated by the observation that the Grid-Minor Theorem can be bypassed in various applications by using Theorems 1.1 and 1.2. We describe two general classes of such applications below.

Bounds for Erdos-Pósa type results. The duality between packing and covering plays a central role in graph theory and combinatorial optimization. One central result of this nature is Menger's theorem, which asserts that for any graph G , subsets S, T of its vertices, and an integer k , either G contains k node-disjoint paths connecting the vertices of S to the vertices of T , or there is a set X of at most $k - 1$ vertices, whose removal disconnects all such paths. Erdos and Pósa [16] proved that for every graph G , either G contains k node-disjoint cycles, or there is a set X of $O(k \log k)$ nodes, whose removal from G makes the graph acyclic. More generally, a family \mathcal{F} of graphs is said to satisfy the Erdos-Pósa property, iff there is an integer-valued function f , such that for every graph G , either G contains k disjoint subgraphs isomorphic to members of \mathcal{F} , or there is a set S of $f(k)$ nodes, such that $G - S$ contains no subgraph isomorphic to a member of \mathcal{F} . In other words, S is a cover, or a hitting set, for \mathcal{F} in G . Erdos-Pósa-type results provide relationships between integral covering and packing problems, and are closely related to fractional covering problems and the integrality gaps of the corresponding LP relaxations.

As an illustrative example for the Erdos-Pósa-type results, let \mathcal{F}_m denote the family of all cycles of length 0 modulo m . Thomassen [35] has proved an Erdos-Pósa-type result for \mathcal{F}_m , by showing that for each graph G , either G contains k disjoint copies of cycles from \mathcal{F}_m , or there is a subset S of $f(k)$ vertices, whose removal disconnects all such cycles in G (here, $f(k) = 2^{2^{O(k)}}$, and m is considered to be a constant).

The proof consists of two steps. In the first step, a simple inductive argument is used to show that for any graph G of treewidth at most w , either G contains k disjoint copies of cycles from \mathcal{F}_m , or there is a subset S of $O(kw)$ vertices, whose removal from G disconnects all such cycles. The second step is to show that if G has treewidth at least some value $g(k)$, then it must contain k disjoint copies of cycles from \mathcal{F}_m . Combining these two steps together, we obtain that $f(k) = O(k \cdot g(k))$. The second step uses Theorem 1.3 to show that, if $\text{tw}(G) \geq g(k) = 2^{m^{O(k)}}$, then G contains a grid minor of size $k(2m)^{2k-1} \times k(2m)^{2k-1}$. This grid minor is then in turn used to find k disjoint copies of cycles from \mathcal{F}_m in G , giving $f(k) = 2^{m^{O(k)}}$.

Using Theorem 1.1, we can significantly strengthen this result, and obtain $f(k) = \tilde{O}(k)$, as follows.³ Assume first that we are given any graph G , with $\text{tw}(G) \geq f'(m)k \text{ poly log } k$, where $f'(m)$ is some function of m . Then, using Theorem 1.1, we can partition G into k vertex-disjoint subgraphs of treewidth at least $f'(m)$ each. Using known techniques (such as, e.g., Theorem 1.3), we can then show that each such subgraph must contain a copy of a cycle from \mathcal{F}_m . Therefore, if $\text{tw}(G) \geq f'(m)k \text{ poly log } k$, then G contains k disjoint copies of cycles from \mathcal{F}_m . Combining this with Step 1 of the algorithm of Thomassen, we conclude that every graph G either contains k copies of cycles from \mathcal{F}_m , or there is a subset S of $f(k) = \tilde{O}(k^2)$ vertices, whose removal from G disconnects all such cycles; a stronger bound of $f(k) = \tilde{O}(k)$ can be obtained by refining the Step 1 argument using a divide and conquer analysis [19].

There is a large body of work in graph theory and combinatorics on Erdos-Pósa-type results. Several of these rely on the Grid-Minor Theorem, and consequently the function $f(k)$ is shown to be exponential (or even worse) in k . Some fundamental results in this area can be improved to obtain a bound polynomial in k , using Theorem 1.1 and the general framework outlined above. For example, Robertson and Seymour [32] derived the following as an important consequence of the Grid-Minor Theorem. Given any fixed graph H , let $\mathcal{F}(H)$ be the family of all graphs that contain H as a minor. Then $\mathcal{F}(H)$ has the Erdos-Pósa property iff H is planar. However, the bound they obtained for $f(k)$ is exponential in k . Using the above general framework, we can show that $f(k) = O(k \cdot \text{poly log}(k))$ for any fixed H .

Improved running times for Fixed-Parameter Tractability. The theory of bidimensionality [10] is a powerful methodology in the design of fixed-parameter tractable (FPT) algorithms. It led to sub-exponential (in the parameter k) time FPT algorithms for bidimensional parameters (formally defined in Section 4) in planar graphs, and more generally graphs that exclude a fixed graph H as a minor. The theory is based on the Grid-Minor Theorem. However, in general graphs, the weak bounds of the Grid-Minor Theorem meant that one could only derive FPT algorithms with running time of the form $2^{2^{O(k^{2.5})}} n^{O(1)}$, as shown by Demaine and Hajiaghayi [14]. Our results lead to algorithms with running times of the form $2^{k \text{ poly log}(k)} n^{O(1)}$ for the same class of problems as in [14]. Thus, one can obtain FPT algorithms for a large class of problems in general graphs via a generic methodology, where the running time has a singly-

³Throughout the paper we use \tilde{O} notation to suppress poly-logarithmic factors.

exponential dependence on the parameter k .

The thrust of this paper is to prove Theorems 1.1 and 1.2, and to highlight their applicability as general tools. The applications described in Section 4 are of that flavor; we have not attempted to examine specific problems in depth. We believe that the theorems, and the technical ideas in their proofs, will have further applications.

Overview of techniques and discussion. A significant contribution of this paper is the formulation of the decomposition theorems for treewidth, and identifying their applications. The main new and non-trivial technical contribution is the proof of Theorem 1.1. The proof of Theorem 1.2 is similar in spirit to the recent work of [8] and [5], who obtained a special case of Theorem 1.2 with $h = \text{poly log } k$, and used it to design algorithms for low-congestion routing in undirected graphs. We note that Theorem 1.1 gives a substantially different tradeoff between the parameters h, r and k , when compared to Theorem 1.2, and leads to the improved results for the two applications we mentioned earlier. Its proof uses new ingredients with a connection to decomposing expanders as explained below.

Contracted graph, well-linked decomposition, and expanders: The three key technical ingredients in the proof of Theorem 1.1 are in the title of the paragraph. To illustrate some key ideas we first consider how one may prove Theorem 1.1 if G is an n -vertex constant-degree expander, which has treewidth $\Omega(n)$. At a high level, one can achieve this as follows. We can take h disjoint copies of an expander with $\Omega(r)$ nodes each (the expansion certifies that treewidth of each copy is r), and “embed” them into G in a vertex-disjoint fashion. This is roughly possible, modulo various non-trivial technical issues, using short-path vertex-disjoint routing in expanders [24]. Now consider a general graph G . For instance it can be a planar graph on n nodes with treewidth $O(\sqrt{n})$; note that the ratio of treewidth to the number of nodes is very different than that in an expander. Here we employ a different strategy, where we cut along a small separator and retain large treewidth on both sides and apply this iteratively until we obtain the desired number of subgraphs. The non-trivial part of the proof is to be able to handle these different scenarios. Another technical difficulty is the following. Treewidth of a graph is a global parameter and there can be portions of the graph that can be removed without changing the treewidth. It is not easy to cleanly characterize the minimal subgraph of G that has roughly the same treewidth as that of G . A key technical ingredient here is borrowed from previous work on graph decompositions [27, 8], namely, the notion of a contracted graph. The contracted graph tries to achieve this minimality, by contracting portions of the graph that satisfy the following technical condition: they have a small boundary and the boundary is well-linked with respect to the contracted portion. Finally, a recurring technical ingredient is the notion of a well-linked decomposition. This allows us to remove a small number of edges while ensuring that the remaining pieces have good conductance. This high-level clustering idea has been crucial in many applications.

Related work on grid-like minors and (perfect) brambles: An important ingredient in the decomposition results is a need to certify that the treewidth of a given graph is large, say at least r . Interestingly, despite being NP-Hard to compute, the treewidth of a graph G has an exact

min-max formula via the bramble number [34] (see Section 2). However, Grohe and Marx [20] have shown that there are graphs G (in fact expanders) for which a polynomial-sized bramble can only certify that treewidth of G is $\Omega(\sqrt{k})$ where $k = \text{tw}(G)$; certifying that G has larger treewidth would require super-polynomial sized brambles. Kreutzer and Tamari [23], building on [20], gave efficient algorithms to construct brambles of order $\tilde{\Omega}(\sqrt{k})$. They also gave efficient algorithms to compute “grid-like” minors introduced by Reed and Wood [30] where it is shown that G has a grid-like minor of size ℓ as long as $\text{tw}(G) = \Omega(\ell^4 \sqrt{\log \ell})$. In some applications it is feasible to use a grid-like minor in place of a grid and obtain improved results. Kreutzer and Tamari [23] used them to define perfect brambles and gave a meta-theorem to obtain FPT algorithms, for a subclass of problems considered in [14], with a single-exponential dependence on the parameter k . Our approach in this paper is different, and in a sense orthogonal, as we explain below. First, a grid-like minor is a single connected structure that does not allow for a decomposition into disjoint grid-like minors. This limitation means one needs a global argument to show that a grid-like minor of a certain size implies a lower bound on some parameter of interest. In contrast, our theorems are specifically tailored to decompose the graph and then apply a local argument in each subgraph, typically to prove that the parameter is at least one in each subgraph. The advantage of our approach is that it is agnostic to how one proves a lower bound in each subgraph; we could use the Grid-Minor Theorem or the more efficient grid-like minor theorem in each subgraph. Kreutzer and Tazari [23] derive efficient FPT algorithms for a subclass of problems considered in [14] where the class is essentially defined as those problems for which one can use a grid-like minor in place of a grid in the global sense described above. In contrast, we can generically handle all the problems considered in [14] as explained in Section 4.

Second, we discuss the efficiency gains possible via our approach. It is well-known that an α -approximation for sparse vertex separators gives an $O(\alpha)$ -approximation for treewidth. Feige et al. [17] obtain an $O(\sqrt{\log \text{tw}(G)})$ -approximation for treewidth. Thus we can efficiently certify treewidth to within a much better factor via separators than with brambles. More explicitly, well-linked sets provide a compact certificate for treewidth; informally, a set of vertices X is well-linked in G if there are no small separators for X — see Section 2 for formal definitions. The tradeoffs we obtain through well-linked sets are stronger than via brambles. In particular, the FPT algorithms that we obtain have a running time $2^{k^{\text{poly log}(k)}} n^{O(1)}$ where k is the parameter of interest. In contrast the algorithms obtained via perfect brambles in [23] have running times of the form $2^{\text{poly}(k)} n^{O(1)}$ where the polynomial is incurred due to the inefficiency in the relationship between treewidth and the size of a grid-like minor. Although the precise dependence on k depends on the parameter of interest, the current bounds require at least a quadratic dependence on k .

Organization: Most of the proofs, and in particular, all the details of proof of Theorem 1.2 are omitted due to space constraints, and can be found in the full version of the paper. Section 3 describes our proof of Theorem 1.1. Section 4 describes the applications and it relies only on the statement of Theorem 1.1, and can be read independently.

2. PRELIMINARIES AND NOTATION

Given a graph G and a set of vertices A , we denote by $\text{out}_G(A)$ the set of edges with exactly one end point in A and by $E_G(A)$ the set of edges with both end points in A . For disjoint sets of vertices A, B the set of edges with one end point in A and the other in B is denoted by $E_G(A, B)$. When clear from context, we omit the subscript G . All logarithms are to the base of 2. We use the following simple claim several times.

Claim 2.1 *Let $\{x_1, \dots, x_n\}$ be a set of non-negative integers, with $\sum_i x_i = N$, and $x_i \leq 2N/3$ for all i . Then we can efficiently compute a partition (A, B) of $\{1, \dots, n\}$, such that $\sum_{i \in A} x_i, \sum_{i \in B} x_i \geq N/3$.*

Graph partitioning. Suppose we are given any graph $G = (V, E)$ with a set T of vertices called terminals. Given any partition (S, \bar{S}) of $V(G)$, the *sparsity* of the cut (S, \bar{S}) with respect to T is $\Phi(S, \bar{S}) = \frac{|E(S, \bar{S})|}{\min\{|T \cap S|, |T \cap \bar{S}|\}}$. The *conductance* of the cut (S, \bar{S}) is $\Psi(S, \bar{S}) = \frac{|E(S, \bar{S})|}{\min\{|E(S)|, |E(\bar{S})|\}}$. We then denote: $\Phi(G) = \min_{S \subset V} \{\Phi(S, \bar{S})\}$, and $\Psi(G) = \min_{S \subset V} \{\Psi(S, \bar{S})\}$. Arora, Rao and Vazirani [3] showed an algorithm that, given a graph G with a set T of k terminals, produces a cut (S, \bar{S}) with $\Phi(S, \bar{S}) \leq \alpha_{\text{ARV}}(k) \cdot \Phi(G)$, where $\alpha_{\text{ARV}}(k) = O(\sqrt{\log k})$. Their algorithm can also be used to find a cut (S, \bar{S}) with $\Psi(S, \bar{S}) \leq \alpha_{\text{ARV}}(m) \cdot \Psi(G)$, where $m = |E(G)|$. We denote this algorithm by \mathcal{A}_{ARV} , and its approximation factor by α_{ARV} from now on.

Well-linkedness and treewidth. The treewidth of a graph $G = (V, E)$ is typically defined via tree decompositions. A tree-decomposition for G consists of a tree $T = (V(T), E(T))$ and a collection of sets $\{X_v \subseteq V\}_{v \in V(T)}$ called bags, such that the following two properties are satisfied: (i) for each edge $ab \in E$, there is some node $v \in V(T)$ with both $a, b \in X_v$ and (ii) for each vertex $a \in V$, the set of all nodes of T whose bags contain a form a connected subtree of T . The *width* of a given tree decomposition is $\max_{v \in V(T)} |X_v| - 1$, and the treewidth of a graph G , denoted by $\text{tw}(G)$, is the width of a minimum-width tree decomposition for G .

It is convenient to work with well-linked sets instead of treewidth. We describe the relationship between them after formally defining the notion of well-linkedness that we require.

Definition 2.1 *We say that a set T of vertices is α -well-linked⁴ in G , iff for any partition (A, B) of the vertices of G into two subsets, $|E(A, B)| \geq \alpha \cdot \min\{|A \cap T|, |B \cap T|\}$.*

Definition 2.2 *We say that a set T of vertices is node-well-linked in G , iff for any pair (T_1, T_2) of equal-sized subsets of T , there is a collection \mathcal{P} of $|T_1|$ **node-disjoint** paths, connecting the vertices of T_1 to the vertices of T_2 . (Note that T_1, T_2 are not necessarily disjoint, and we allow empty paths).*

⁴This notion of well-linkedness is based on edge-cuts and we distinguish it from node-well-linkedness that is directly related to treewidth. For technical reasons it is easier to work with edge-cuts and hence we use the term well-linked to mean edge-well-linkedness, and explicitly use the term node-well-linkedness when necessary.

Lemma 2.1 (Reed [29]) *Let k be the size of the largest node-well-linked set in G . Then $k \leq \text{tw}(G) \leq 4k$.*

We then obtain the following simple corollary, whose proof appears in the full version.

Corollary 2.1 *Let G be any graph with maximum vertex degree at most Δ , and let T be any subset of vertices, such that T is α -well-linked in G , for some $0 < \alpha < 1$. Then $\text{tw}(G) \geq \frac{\alpha \cdot |T|}{3\Delta} - 1$.*

Lemma 2.1 guarantees that any graph G of treewidth k contains a set X of $\Omega(k)$ vertices, that is vertex-well-linked in G . Kreutzer and Tazari [23] give a constructive version of this lemma, obtaining a set X with slightly weaker properties. Below is a rephrasing of Lemma 3.7 in [23] in terms convenient to us.

Lemma 2.2 *There is a polynomial-time algorithm, that, given a graph G of treewidth k , finds a set X of $\Omega(k)$ vertices, such that X is $\alpha^* = \Omega(1/\log k)$ -well-linked in G . Moreover, for any partition (X_1, X_2) of X into two equal-sized subsets, there is a collection \mathcal{P} of paths connecting every vertex of X_1 to a distinct vertex of X_2 , such that every vertex of G participates in at most $1/\alpha^*$ paths in \mathcal{P} .*

Well-linked decompositions. Let S be any subset of vertices in G . Informally, we say that S is α -good iff the subset $|\text{out}(S)|$ of edges is α -well-linked in the graph $G[S] \cup \text{out}(S)$. Formally, S is α -good⁵, iff for any partition (A, B) of S , $|E(A, B)| \geq \alpha \cdot \min\{|\text{out}(A) \cap \text{out}(S)|, |\text{out}(B) \cap \text{out}(S)|\}$. A set $D : \text{out}(S) \times \text{out}(S) \rightarrow \mathbb{R}^+$ of demands defines, for every pair $e, e' \in \text{out}(S)$, a demand $D(e, e')$. We say that D is a c -restricted set of demands, iff for every $e \in \text{out}(S)$, $\sum_{e' \in \text{out}(S)} D(e, e') \leq c$. Assume that S is an α -good subset of vertices in G . From the duality of cuts and flows, and from the known bounds on the flow-cut gap in undirected graphs [25], if D is any set of c -restricted demands over $\text{out}(S)$, then it can be fractionally routed inside $G[S]$ with edge-congestion at most $O(c \log k'/\alpha)$, where $k' = |\text{out}(S)|$. The following theorem, in its many variations, (sometimes under the name of "well-linked decomposition") has been used extensively in routing and graph decomposition (see e.g. [27, 6, 7, 28, 2, 8, 9, 5]). The proof appears in the full version.

Theorem 2.1 *Let S be any subset of vertices of G , with $|\text{out}(S)| = k'$, and let $0 < \alpha < \frac{1}{8\alpha_{\text{ARV}}(k') \cdot \log k'}$ be a parameter. Then there is an efficient algorithm to compute a partition \mathcal{W} of S , such that for each $W \in \mathcal{W}$, $|\text{out}(W)| \leq k'$ and W is α -good. Moreover, $\sum_{W \in \mathcal{W}} |\text{out}(W)| \leq k'(1 + 16\alpha \cdot \alpha_{\text{ARV}}(k') \cdot \log k') = k'(1 + O(\alpha \log^{3/2} k'))$. The parameter $\alpha_{\text{ARV}}(k')$ can be set to 1 if the efficiency of the algorithm is not relevant.*

Pre-processing to reduce maximum degree. Let G be any graph with $\text{tw}(G) = k$. The proofs of Theorems 1.1 and 1.2 work with edge-well-linked sets instead of the node-well-linked ones. In order to be able to translate between both types of well-linkedness and the treewidth, we need to reduce

⁵The same property was called "bandwidth property" in [27], and in [8, 9], set S with this property was called α -well-linked. We choose this notation to avoid confusion with other notions of well-linkedness used in this paper.

the maximum vertex degree of the input graph G . Using the cut-matching game of Khandekar, Rao and Vazirani [22], one can reduce the maximum vertex degree to $O(\log^3 k)$, while only losing a poly $\log k$ factor in the treewidth, as was noted in [5] (see Remark 2.2). We state the theorem formally below. A brief proof sketch appears in the full version.

Theorem 2.2 *Let G be any graph with treewidth k . Then there is an efficient randomized algorithm to compute a subgraph G' of G , with maximum vertex degree at most $O(\log^3 k)$ such that $\text{tw}(G') = \Omega(k/\log^6 k)$.*

Remark 2.3 *In fact a stronger result, giving a constant bound on the maximum degree follows from the expander embedding result in [5]. However, the bound on the treewidth guaranteed is worse than in the preceding theorem by a (large) polylogarithmic factor. For our algorithms, the polylogarithmic bound on the degree guaranteed by Theorem 2.2 is sufficient.*

3. PROOF OF THEOREM 1.1

We start with a graph G whose treewidth is at least k . For our algorithm, we need to know the value of the treewidth of G , instead of the lower bound on it. We can compute the treewidth of G approximately, to within an $O(\log(\text{tw}(G)))$ -factor, using the algorithm of Amir [1]. Therefore, we assume that we are given a value $k' \geq k$, such that $\Omega(k'/\log k') \leq \text{tw}(G) \leq k'$.

We then apply Theorem 2.2, to obtain a subgraph G' of G of maximum vertex degree $\Delta = O(\log^3 k')$ and treewidth $\Omega(k'/\log^7 k')$. Using Lemma 2.2, we compute a subset T of $\Omega(k'/\log^7 k')$ vertices, such that T is $\Omega(1/\log k')$ -well-linked in G' .

In order to simplify the notation, we denote G' by G and $|T|$ by k from now on. From the above discussion, $\text{tw}(G) \leq ck \log^7 k$ for some constant c , T is $\Omega(1/\log k)$ -well-linked in G , and the maximum vertex degree in G is $\Delta = O(\log^3 k)$; we define the parameter α^* to be $\Omega(1/\log k)$ which is the well-linkedness guarantee given by Lemma 2.2. It is now enough to find a collection G_1, \dots, G_h of vertex-disjoint subgraphs of G , such that $\text{tw}(G_i) \geq r$ for each i . We use the parameter $r' = c' \Delta^2 r \log^{11} k$, where c' is a sufficiently large constant. We assume without loss of generality that k is large enough, so, for example, $k \geq c'' r \log^{30} k$, where c'' is a large enough constant. We also assume without loss of generality that G is connected.

Definition 3.1 *We say that a subset S of vertices in G is an acceptable cluster, iff $|\text{out}(S)| \leq r'$, $|S \cap T| \leq |T|/2$, and S is α_G -good, for $\alpha_G = \frac{1}{256\alpha_{\text{ARV}}(k) \log k} = \Theta\left(\frac{1}{\log^{1.5} k}\right)$.*

Notice that since the maximum vertex degree in G is bounded by $\Delta < r'$, if S consists of a single vertex, then it is an acceptable cluster. Given any partition \mathcal{C} of the vertices of G into acceptable clusters, we let $H_{\mathcal{C}}$ be the *contracted graph* associated with \mathcal{C} . Graph $H_{\mathcal{C}}$ is obtained from G by contracting every cluster $C \in \mathcal{C}$ into a single vertex v_C , that we refer to as a super-node. We delete self-loops, but leave parallel edges. Notice that the maximum vertex degree in $H_{\mathcal{C}}$ is bounded by r' . We denote by $\varphi(\mathcal{C})$ the total number of edges in $H_{\mathcal{C}}$. Below is a simple observation that follows

from the α^* -well-linkedness of T in G . The proof appears in the full version of the paper.

Observation 3.1 *Let \mathcal{C} be any partition of the vertices of G into acceptable clusters. Then $\varphi(\mathcal{C}) \geq \alpha^* k/3$.*

Throughout the algorithm, we maintain a partition \mathcal{C} of $V(G)$ into acceptable clusters. At the beginning, $\mathcal{C} = \{\{v\} \mid v \in V(G)\}$. We then perform a number of iterations. In each iteration, we either compute a partition of G into h disjoint sub-graphs, of treewidth at least r each, or find a new partition \mathcal{C}' of $V(G)$ into acceptable clusters, such that $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$. The execution of each iteration is summarized in the following theorem.

Theorem 3.1 *There is an efficient algorithm, that, given a partition \mathcal{C} of $V(G)$ into acceptable clusters, either computes a partition of G into h disjoint subgraphs of treewidth at least r each, or returns a new partition \mathcal{C}' of $V(G)$ into acceptable clusters, such that $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$.*

Clearly, after applying Theorem 3.1 at most $|E(G)|$ times, we obtain a partition of G into h disjoint subgraphs of treewidth at least r each. From now on we focus on proving Theorem 3.1. Given a current partition \mathcal{C} of $V(G)$ into acceptable clusters, let H denote the corresponding contracted graph. We denote $n = |V(H)|$, $m = |E(H)|$. Notice that from Observation 3.1, $m \geq \alpha^* k/3$. We now consider two cases, and prove Theorem 1.1 separately for each of them. The first case is when $n \geq k^5$.

3.1 Case 1: $n \geq k^5$

We note that n is large when compared to the treewidth and hence we expect the graph H should have low expansion. Otherwise, we get a contradiction by showing that $\text{tw}(G) > ck \log^7 k$. The proof strategy in the low-expansion regime is to repeatedly decompose along balanced partitions to obtain h subgraphs with treewidth at least r each.

Let $z = k^5$. The algorithm first chooses an arbitrary subset Z of z vertices from H . Suppose we are given any subset S of vertices of H . We say that a partition (A, B) of S is γ -balanced (with respect to Z), iff $\min\{|A \cap Z|, |B \cap Z|\} \geq \gamma |S \cap Z|$. We say that it is *balanced* iff it is γ -balanced for $\gamma = \frac{1}{4}$. The following claim is central to the proof of the theorem in Case 1.

Claim 3.1 *Let S be any subset of vertices in H with $|S \cap Z| > 100$, and let (A, B) a balanced partition of S (with respect to Z), minimizing $|E_H(A, B)|$. Then $|E_H(A, B)| \leq k^2$.*

The proof shows that if $|E_H(A, B)| > k^2$, then $\text{tw}(G) > ck \log^7 k$, leading to a contradiction. We defer the proof to the full version.

We now show an algorithm to find the desired collection G_1, \dots, G_h of subgraphs of G . We use the algorithm \mathcal{A}_{ARV} of Arora, Rao and Vazirani [3] to find a balanced partition of a given set S of vertices of H ; the algorithm is applied to G with $S \cap Z$ as the terminals. Given any such set S of vertices, the algorithm \mathcal{A}_{ARV} returns a γ_{ARV} -balanced partition (A, B) of S , with $|E_H(A, B)| \leq \alpha_{\text{ARV}}(z) \cdot \text{OPT}$, where OPT is the smallest number of edges in any balanced partition, and γ_{ARV} is some constant. In particular, from Claim 3.1, $|E(A, B)| \leq \alpha_{\text{ARV}}(z) \cdot k^2$, if $|S \cap Z| \geq 100$.

We start with $\mathcal{S} = \{V(H)\}$, and perform h iterations. At the beginning of iteration i , set \mathcal{S} will contain i disjoint subsets of vertices of H . An iteration is executed as follows. We select a set $S \in \mathcal{S}$, maximizing $|Z \cap S|$, and compute a γ_{ARV} -balanced partition (A, B) of S , using the algorithm \mathcal{A}_{ARV} . We then remove S from \mathcal{S} , and add A and B to it instead. Let $\mathcal{S} = \{X_1, \dots, X_{h+1}\}$ be the final collection of sets after h iterations. From Claim 3.1, the increase in $\sum_{X \in \mathcal{S}} |\text{out}_H(X)|$ is bounded by $k^2 \alpha_{\text{ARV}}(z)$ in each iteration. Therefore, throughout the algorithm, $\sum_{X \in \mathcal{S}} |\text{out}_H(X)| \leq k^2 \alpha_{\text{ARV}}(z) h$ holds. In the following observation, whose proof appears in the full version, we show that for each $X_i \in \mathcal{S}$, $|X_i \cap Z| \geq \frac{\gamma_{\text{ARV}} \cdot z}{2h}$.

Observation 3.2 *Consider some iteration i of the algorithm. Let \mathcal{S}_i be the collection of vertex subsets at the beginning of iteration i , let $S \in \mathcal{S}_i$ be the set that was selected in this iteration, and let \mathcal{S}_{i+1} be the set obtained after replacing S with A and B . Then $|A \cap Z|, |B \cap Z| \geq \gamma_{\text{ARV}} \cdot |S \cap Z|$, and for each $S' \in \mathcal{S}_{i+1}$, $|S' \cap Z| \geq \frac{\gamma_{\text{ARV}} \cdot z}{2h}$.*

Among the sets X_1, \dots, X_{h+1} , there can be at most one set X_i , with $|T \cap (\bigcup_{v \in C} X_i)| > |T|/2$. We assume without loss of generality that this set is X_{h+1} , and we will ignore it from now on. Consider now some set X_i , for $1 \leq i \leq h$. Since graph H is connected, and X_i contains at least $\frac{\gamma_{\text{ARV}} \cdot z}{2h}$ vertices (the vertices of $X_i \cap Z$), while $|\text{out}_H(X_i)| \leq k^2 h \alpha_{\text{ARV}}(z)$, it follows that $|E_H(X_i)| \geq \frac{1}{2} (\frac{\gamma_{\text{ARV}} \cdot z}{2h} - k^2 h \alpha_{\text{ARV}}(z)) \geq \frac{\gamma_{\text{ARV}} \cdot z}{8h} > 64 |\text{out}_H(X_i)|$, as $z = k^5$, and k is large enough. Let X'_i be the subset of vertices obtained from X_i , by uncontracting all super-nodes of X_i . Then $|E_G(X'_i)| \geq |E_H(X_i)| \geq 64 |\text{out}_H(X_i)| = 64 |\text{out}_G(X'_i)|$.

Our next step is to compute a decomposition \mathcal{W}_i of X'_i into α_G -good clusters, using Theorem 2.1. Notice that $k' = |\text{out}_G(X'_i)| \leq k^2 h \alpha_{\text{ARV}}(z) \leq 5k^2 h \alpha_{\text{ARV}}(k) < k^4$ since $z = k^5$; therefore the choice of $\alpha_G = \frac{1}{256 \alpha_{\text{ARV}} k \log k} < \frac{1}{8 \alpha_{\text{ARV}}(k') \log k'}$ satisfies the conditions of the theorem.

Assume first that for every cluster $C_i \in \mathcal{W}_i$, $|\text{out}_G(C_i)| \leq r'$. Then we can obtain a new partition \mathcal{C}' of $V(G)$ into acceptable clusters with $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$, as follows. We add to \mathcal{C}' all clusters $C \in \mathcal{C}$ that are disjoint from X'_i , and we add all clusters in \mathcal{W}_i to it as well. Clearly, the resulting partition \mathcal{C}' consists of acceptable clusters only. We now show that $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$. Indeed,

$$\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - |\text{out}_H(X_i)| - |E_H(X_i)| + \sum_{R \in \mathcal{W}_i} |\text{out}_G(R)|$$

From the choice of α_G , $\sum_{R \in \mathcal{W}_i} |\text{out}_G(R)| < 3 |\text{out}_G(X'_i)| = 3 |\text{out}_H(X_i)|$ holds, while $|E_H(X_i)| \geq 64 |\text{out}_H(X_i)|$. Therefore, $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$.

Assume now that for each $1 \leq i \leq h$, there is at least one cluster $C_i \in \mathcal{W}_i$ with $|\text{out}_G(C_i)| \geq r'$. Let $\{C_1, \dots, C_h\}$ be the resulting collection of clusters, where for each i , $C_i \in \mathcal{W}_i$. For $1 \leq i \leq h$, we now let $G_i = G[C_i]$. It is easy to see that the graphs G_1, \dots, G_h are vertex-disjoint. It now only remains to show that each graph G_i has treewidth at least r . Fix some $1 \leq i \leq h$, and let $\Gamma_i \subseteq C_i$ contain the endpoints of edges in $\text{out}_G(C_i)$, that is, $\Gamma_i = \{v \in C_i \mid \exists e = (u, v) \in \text{out}_G(C_i)\}$. Then, since C_i is an α_G -good set of vertices, Γ_i is α_G -well-linked in the graph G_i . Moreover,

$$|\Gamma_i| \geq |\text{out}(C_i)|/\Delta \geq r'/\Delta. \text{ From Corollary 2.1, } \text{tw}(G_i) \geq \frac{\alpha_G r'}{3\Delta^2} - 1 \geq r.$$

3.2 Case 2: $n < k^5$

Since vertex degrees in H are bounded by r' , $m = O(k^5 r') = O(k^6)$. The algorithm for Case 2 consists of two phases. In the first phase, we partition $V(H)$ into a number of disjoint subsets X_1, \dots, X_ℓ , where, on the one hand, for each X_i , the conductance of $H[X_i]$ is large, while, on the other hand, $\sum_{i=1}^\ell |\text{out}(X_i)| \leq |E(H)|/10$. We discard all clusters X_i with $|\text{out}(X_i)| \geq |E(X_i)|/2$, denoting by \mathcal{X} the collection of the remaining clusters, and show that $\sum_{X_i \in \mathcal{X}} |E(X_i)| = \Omega(\alpha^* k)$. If any cluster $X \in \mathcal{X}$ has $|E(X)| \leq 2r'$, then we find a new partition \mathcal{C}' of the vertices of G into acceptable clusters, with $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$. Therefore, we can assume that for every cluster $X \in \mathcal{X}$, $|E(X)| > 2r'$. We then proceed to the second phase. Here, we take advantage of the high conductance of each $X_i \in \mathcal{X}$ to show that X_i can be partitioned into h_i vertex-disjoint sub-graphs, such that we can embed a large enough expander into each such subgraph. The value h_i is proportional to $|E(X_i)|$, and we ensure that $\sum_{X_i \in \mathcal{X}} h_i \geq h$ to get the desired number of subgraphs. The embedding of the expander into each sub-graph is then used as a certificate that this sub-graph (or more precisely, a sub-graph of G obtained after uncontracting the super-nodes) has large treewidth.

Phase 1.

We use the following theorem, that allows us to decompose any graph into a collection of high-conductance connected components, by only removing a small fraction of the edges. A similar procedure has been used in previous work, and can be proved using standard graph decomposition techniques. The proof is deferred to the full version.

Theorem 3.2 *Let H be any connected n -vertex graph containing m edges. Then there is an efficient algorithm to compute a partition X_1, \dots, X_ℓ of the vertices of H , such that: (i) for each $1 \leq i \leq \ell$, the conductance of graph $H[X_i]$, $\Psi(H[X_i]) \geq \frac{1}{160 \alpha_{\text{ARV}}(m) \log m}$; and (ii) $\sum_{i=1}^\ell |\text{out}(X_i)| \leq m/10$.*

The algorithm in phase 1 uses Theorem 3.2 to partition the contracted graph H into a collection $\{X_1, \dots, X_\ell\}$ of clusters. Recall that $m = |E(H)|$, and $n = |V(H)|$. We are guaranteed that $\sum_{i=1}^\ell |E(X_i)| \geq 0.9m$ and $\sum_{i=1}^\ell |\text{out}(X_i)| \leq 0.1m$ from Theorem 3.2.

Let \mathcal{X}' contain all clusters X_i with $|\text{out}(X_i)| \geq \frac{1}{2}|E(X_i)|$, and let \mathcal{X} contain all remaining clusters. Notice that $\sum_{X_i \in \mathcal{X}'} |E(X_i)| \leq 2 \sum_{X_i \in \mathcal{X}'} |\text{out}(X_i)| \leq 2 \sum_{i=1}^\ell |\text{out}(X_i)|$, which is at most $0.2m$. Therefore, $\sum_{X_i \in \mathcal{X}} |E(X_i)| \geq \frac{1}{2}m \geq \frac{\alpha^* k}{6}$ from Observation 3.1. From now on we only focus on clusters in \mathcal{X} .

Assume first that there is some cluster $X_i \in \mathcal{X}$, with $|E(X_i)| \leq 2r'$. We claim that in this case, we can find a new partition \mathcal{C}' of the vertices of G into acceptable clusters, with $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$. We first need the following simple observation whose proof appears in the full version.

Observation 3.3 *Assume that for some $X_i \in \mathcal{X}$, $|E(X_i)| \leq 2r'$ holds. Let $X' = \bigcup_{v \in C} X_i$. Then $|X' \cap T| < |T|/2$.*

Let X'_i be the set of vertices of G , obtained from X_i by uncontracting the super-nodes of X_i . We apply Theorem 2.1 to the set X'_i of vertices, to obtain a partition \mathcal{W}_i of X'_i into α_G -good clusters. It is easy to see that all clusters in \mathcal{W}_i are acceptable, since we are guaranteed that for each $R \in \mathcal{W}_i$, $|\text{out}(R)| \leq |\text{out}(X'_i)| \leq r'$, and $|R \cap T| < |T|/2$. The new partition \mathcal{C}' of the vertices of G into acceptable clusters is obtained as follows. We include all clusters of \mathcal{C} that are disjoint from X'_i , and we additionally include all clusters in \mathcal{W}_i . From the above discussion, all clusters in \mathcal{C}' are acceptable. It now only remains to bound $\varphi(\mathcal{C}')$. It is easy to see that $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - |E_H(X_i)| - |\text{out}_H(X_i)| + \sum_{R \in \mathcal{W}_i} |\text{out}(R)|$. The choice of α_G ensures that $\sum_{R \in \mathcal{W}_i} |\text{out}(R)| \leq 1.25|\text{out}_G(X'_i)| = 1.25|\text{out}_H(X_i)|$. Since $|E_H(X_i)| > 2|\text{out}_H(X_i)|$, we get that $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$. From now on, we assume that for every cluster $X_i \in \mathcal{X}$, $|E(X_i)| \geq 2r'$.

Phase 2.

For convenience, we assume without loss of generality, that $\mathcal{X} = \{X_1, \dots, X_z\}$. For $1 \leq i \leq z$, let $m_i = |E(X_i)|$. Recall that from the above discussion, $m_i \geq r'$, and $\sum_{i=1}^z m_i \geq \alpha^* k/6$. We set $h_i = \lceil \frac{6m_i h}{\alpha^* k} \rceil$. Let $X'_i = \bigcup_{v \in C \in X_i} C$. In the remainder of this section, we will partition, for each $1 \leq i \leq z$, the graph $G[X'_i]$ into h_i vertex-disjoint subgraphs, of treewidth at least r each. Since $\sum_{i=1}^z h_i \geq \sum_{i=1}^z \frac{6m_i h}{\alpha^* k} \geq h$, this will complete the proof of Theorem 1.1.

From now on, we focus on a specific graph $H[X_i]$, and its corresponding uncontracted graph $G[X'_i]$. Our algorithm performs h_i iterations. In the first iteration, we embed an expander over $r'' = r \text{ poly log } k$ vertices into $H[X_i]$. We then partition $H[X_i]$ into two sub-graphs: H_1 , containing all vertices that participate in this embedding, and H'_1 containing all remaining vertices. Our embedding will ensure that $\sum_{v \in V(H_1)} d_H(v) \leq r^2 \text{ poly log } k$, or in other words, we can obtain H'_1 from $H[X_i]$ by removing only $r^2 \text{ poly log } k$ edges from it, and deleting isolated vertices. We then proceed to the second iteration, and embed another expander over r'' vertices into H'_1 . This in turn partitions H'_1 into H_2 , that contains the embedding of the expander, and H'_2 , containing the remaining edges. In general, in iteration j , we start with a sub-graph H'_{j-1} of H , and partition it into two subgraphs: H_j containing the embedding of an expander, and H'_j that becomes an input to the next iteration. Since we ensure that for each graph H_j , the total out-degree of its vertices is bounded by $r^2 \text{ poly log } k$, each residual graph H'_j is guaranteed to contain a large fraction of the edges of the original graph $H[X_i]$. We show that this in turn guarantees that H'_j contains a large sub-graph with a large conductance, which will in turn allow us to embed an expander over a subset of r'' vertices into H'_j in the following iteration.

We start with the following theorem that forms the technical basis for iteratively embedding multiple expanders of certain size into a larger expander. The proof appears in the full version.

Theorem 3.3 *Let \mathbf{G} be any graph with $|E(\mathbf{G})| = m$ and $\Psi(\mathbf{G}) \geq \gamma$, where $\gamma \leq 0.1$. Let \mathbf{H} be a sub-graph obtained from \mathbf{G} by removing some subset S_0 of vertices and all their adjacent edges, so $\mathbf{H} = \mathbf{G} - S_0$. Assume further that $|E(\mathbf{G}) \setminus E(\mathbf{H})| \leq \gamma m/8$. Then we can efficiently compute a subset S of vertices in \mathbf{H} , such that $\mathbf{H}[S]$ contains at least $m/2$ edges*

and has conductance at least $\frac{\gamma}{4\alpha_{\text{ARV}}(m)}$.

The next theorem is central to the execution of Phase 2. The theorem shows that, if we are given a sub-graph H' of H that has a high enough conductance, and contains at least r' edges, then we can find a subset S of r' vertices of H' , such that the following holds: if $S' = \bigcup_{v \in C \in S} C$, and $G' = G[S']$, then $\text{tw}(G') \geq r$. In order to show this, we embed an expander over a set of $r'' = r \text{ poly log } k$ of vertices into H' , and define S to be the set of all vertices of H' participating in this embedding. The embedding of the expander into $H'[S]$ is then used to certify that the treewidth of the resulting graph G' is at least r . The proof of the following theorem appears in the full version.

Theorem 3.4 *Let H' be any vertex-induced subgraph of H , such that $|E(H')| \geq r'$, and $\Psi(H') \geq \frac{1}{640\alpha_{\text{ARV}}^2(m) \log m}$. Then there is an efficient algorithm to find a subset S of at most r' vertices of H' , such that, if G' is obtained from $H'[S]$ by uncontracting the super-nodes in S , then $\text{tw}(G') \geq r$.*

We are now ready to complete the description of the algorithm for Phase 2. Our algorithm considers each one of the subsets $X_i \in \mathcal{X}$ of vertices separately. Fix some $1 \leq i \leq z$. If $h_i = 1$, then by Theorem 3.4, graph $G[X_i]$ has treewidth at least r . Otherwise, we perform h_i iterations. At the beginning of every iteration j , we are given some vertex-induced subgraph H_j of $H[X_i]$, with $|E(H_j)| \geq m_i/2$ and $\Psi(H_j) \geq \frac{1}{640\alpha_{\text{ARV}}^2(m) \log m}$. At the beginning, $H_1 = H[X_i]$, and as observed before, $\Psi(H_1) \geq \frac{1}{160\alpha_{\text{ARV}}(m) \log m} \geq \frac{1}{640\alpha_{\text{ARV}}^2(m) \log m}$. In order to execute the j th iteration, we apply Theorem 3.4 to graph $H' = H_j$, and compute a subset S of at most r' vertices of H' . We denote this set of vertices by S_j^i , and we let $H_j^i = H[S_j^i]$. We also let G_j^i be the sub-graph of G , obtained by uncontracting the super-nodes of H_j^i . From Theorem 3.4, $\text{tw}(G_j^i) \geq r$. We then apply Theorem 3.3 to graph $\mathbf{G} = H[X_i]$, set $S_0 = \bigcup_{j'=1}^j S_{j'}^i$, and $\mathbf{H} = \mathbf{G} \setminus S_0$, to obtain the graph $H_{j+1} = \mathbf{H}[S]$ that becomes an input to the next iteration.

In order to show that we can carry this process out for h_i iterations, it is enough to prove that $\sum_{j=1}^{h_i} \sum_{v \in S_j^i} d_H(v) \leq \gamma m_i/8$, where $\gamma = \frac{1}{160\alpha_{\text{ARV}}(m) \log m}$. Indeed, since the vertex degrees in H are bounded by r' ,

$$\sum_{j=1}^{h_i} \sum_{v \in S_j^i} d_H(v) \leq \sum_{j=1}^{h_i} r' \cdot |S_j^i| \leq (r')^2 \cdot h_i \leq O\left(\frac{m_i r^2 h \text{ poly log } k}{\alpha^* k}\right),$$

by substituting $h_i = \lceil \frac{6m_i h}{\alpha^* k} \rceil$. Since we assume that $r^2 h < k/\text{poly log } k$, and $m = O(k^6)$, it follows that the sum is bounded by $\frac{m_i}{1280\alpha_{\text{ARV}}(m) \log m}$, as required.

Our final collection of subgraphs is $\Pi = \{G_j^i \mid 1 \leq i \leq z, 1 \leq j \leq h_i\}$. From the above discussion, Π contains $\sum_{i=1}^z h_i \geq h$ subgraphs of treewidth at least r each.

4. APPLICATIONS

We now describe two applications of Theorem 1.1. Consider an integer-valued parameter P that associates a number $P(G)$ with each graph G . For instance, $P(G)$ could be the

size of the smallest vertex cover of G , or it could be the maximum number of vertex-disjoint cycles in G . We say that P is *minor-closed* if $P(G) \geq P(H)$ for any minor H of G , that is, the value does not increase when deleting edges or contracting edges. A number of interesting parameters are minor-closed. Following [10], we say that P has the *parameter-treewidth bound*, if there is some function $f : \mathbb{Z}_+ \rightarrow \mathbb{Z}_+$ such that $P(G) \leq k$ implies that $\text{tw}(G) \leq f(k)$. In other words, if the treewidth of G is large, then $P(G)$ must also be large. A minor-closed property P has the parameter-treewidth bound iff it has the bound on the family of grids. This is because an $r \times r$ grid has treewidth r , and the Grid-Minor Theorem shows that sufficiently large treewidth implies the existence of a large grid minor. This approach also has the advantage that grids are simple and concrete graphs to reason about. However, this approach for proving parameter-treewidth bounds suffers from the (current) quantitative weakness in the Grid-Minor theorem. For a given parameter P , one can of course focus on methods that are tailored to it. Alternatively, good results can be obtained in special classes of graphs such as planar graphs, and graphs that exclude a fixed graph as a minor, due to the linear relationship between the treewidth and the grid-minor size in such graphs. Theorem 1.1 allows for a generic method to change the dependence $f(k)$ from exponential to polynomial, under some mild restrictions. The following subsections describe these applications.

4.1 FPT Algorithms in General Graphs

Let P be any minor-closed graph parameter, and consider the decision problem associated with P : Given a graph G and an integer k , is $P(G) \leq k$? We say that parameter P is *fixed-parameter tractable*, iff there is an algorithm for this decision problem, whose running time is $h(k) \cdot n^{O(1)}$ where n is the size of G and h is a function that depends only on k . There is a vast literature on Fixed-Parameter Tractability (FPT), and we refer the reader to [15, 26, 18, 4]. Observe that for any minor-closed parameter P , and any fixed integer k , the family $\mathcal{F} = \{G \mid P(G) \leq k\}$ of graphs is a minor-closed family. That is, if $G \in \mathcal{F}$, and G' is a minor of G , then $G' \in \mathcal{F}$. Therefore, from the work of Robertson and Seymour on graph minors and the proof of Wagner's conjecture, there is a finite family $H_{\mathcal{F}}$ of graphs, such that \mathcal{F} is precisely the set of all graphs that do not contain any graph from $H_{\mathcal{F}}$ as a minor. In particular, in order to test whether $P(G) \leq k$, we only need to check whether G contains a graph from $H_{\mathcal{F}}$ as a minor, and this can be done in time $O(n^3)$ (where we assume that k is a constant), using the work of Robertson and Seymour. However, even though the family $H_{\mathcal{F}}$ of graphs is known to exist, no explicit algorithms for constructing it are known. The family $H_{\mathcal{F}}$ of course depends on P , and moreover, even for a fixed property P , it varies with the parameter k . Therefore, the theory only guarantees the existence of a non-uniform FPT algorithm for every minor-closed parameter P . For this reason, it is natural to consider various restricted classes of minor-closed parameters. Motivated by the existence of sub-exponential time algorithms on planar and H-minor-free graphs, a substantial line of work has focused on *bidimensional* parameters — see Demaine et al. [13], and the survey in [10]. Demaine and Hajiaghayi [14] proved the following generic theorem on Fixed-Parameter Tractability of minor-closed bidimensional properties that satisfy some mild additional conditions.

Theorem 4.1 ([14]) *Consider a minor-closed parameter P that is positive on some $g \times g$ grid, is at least the sum over the connected components of a disconnected graph, and can be computed in $h(w)n^{O(1)}$ time given a width- w tree decomposition of the graph. Then there is an algorithm that decides whether P is at most k on a graph with n vertices in $[2^{2^{O(g\sqrt{k})^5}} + h(2^{O(g\sqrt{k})^5})]n^{O(1)}$ time.*

The main advantage of the above theorem is its generality. However, its proof uses the Grid-Minor Theorem, and hence the running time of the algorithm is doubly exponential in the parameter k . Demaine and Hajiaghayi also observed, in the following theorem, that the running time can be reduced to singly-exponential in k if the Grid-Minor Theorem can be improved substantially.

Theorem 4.2 ([14]) *Assume that every graph of treewidth greater than $\Theta(g^2 \log g)$ has a $g \times g$ grid as a minor. Then for every minor-closed parameter P satisfying the conditions of Theorem 4.1, there is an algorithm that decides whether $P(G) \leq k$ on any n -vertex graph G in $[2^{O(g^2 k \log(gk))} + h(O(g^2 k \log(gk)))]n^{O(1)}$ time.*

We show below that, via Theorem 1.1, we can bypass the need to improve the Grid-Minor Theorem.

Theorem 4.3 *Consider a minor-closed parameter P that is positive on all graphs with treewidth $\geq p$, is at least the sum over the connected components of a disconnected graph, and can be computed in $h(w)n^{O(1)}$ time given a width- w tree decomposition of the graph. Then there is an algorithm that decides whether P is at most k on a graph with n vertices in $[2^{\tilde{O}(p^2 k)} + h(\tilde{O}(p^2 k))]n^{O(1)}$ time.*

PROOF. Let $k' = \tilde{\Theta}(p^2 k)$. If the given graph G has treewidth greater than k' , then by Theorem 1.1 it can be partitioned into k node-disjoint subgraphs G_1, \dots, G_k where $\text{tw}(G_i) \geq p$ for each i . Let G' be obtained by the union of these disconnected graphs (equivalently we remove the edges that do not participate in the graphs G_i from G). From the assumptions on P , $P(G_i) \geq 1$ for each i , and $P(G') \geq \sum_i P(G_i) \geq k$. Moreover, since P is minor-closed, $P(G) \geq P(G')$. Therefore, if $\text{tw}(G) \geq k' = \tilde{\Omega}(p^2 k)$ then $P(G) \geq k$ must hold.

We use known algorithms, for instance [1], that, given a graph G , either produce a tree decomposition of width at most $4w$ or certify that $\text{tw}(G) > w$ in $2^{O(w)}n^{O(1)}$ time. Using such an algorithm we can detect in $2^{O(k')}n^{O(1)}$ time whether G has treewidth at least k' , or find a tree decomposition of width at most $4k'$.

If $\text{tw}(G) \geq k'$, then, as we have argued above, $P(G) \geq k$. We then terminate the algorithm with a positive answer. Otherwise, $\text{tw}(G) < 4k'$ and we can use the promised algorithm that runs in time $h(4k') \cdot n^{O(1)}$ to decide whether $P(G) < k$ or not. The overall running time of the algorithm is easily seen to be the claimed bound. \square

Remark 4.1 *In the proof of Theorem 4.3 the assumption on P being minor-closed is used only in arguing that $P(G') \geq P(G)$. Thus, it suffices to assume that the parameter P does not increase under edge deletions (in addition to the assumption on P over disconnected components of a graph).*

Note that the running time is singly-exponential in p and k . How does one prove an upper bound on p , the minimum treewidth guaranteed to ensure that the parameter value is positive? For some problems it may be easy to directly obtain a good bound on p . The following corollary shows that one can always use grid minors to obtain a bound on p . The run-time dependence on the grid size g is doubly exponential since we are using the Grid-Minor Theorem, but it is only singly-exponential in the parameter k . Thus, if g is considered to be a fixed constant, we obtain singly-exponential Fixed-Parameter Tractability algorithms in general graphs for all the problems that satisfy the conditions in Theorem 4.1.

Corollary 4.1 *Consider a minor-closed parameter P that is positive on some $g \times g$ grid, is at least the sum over the connected components of a disconnected graph, and can be computed in $h(w)n^{O(1)}$ time given a width- w tree decomposition of the graph. Then there is an algorithm that decides whether $P(G) \leq k$ on a graph G with n vertices in $[2^{\tilde{O}(k \cdot 2^{O(g^5)})} + h(\tilde{O}(2^{O(g^5)}k))]n^{O(1)}$ time.*

PROOF. If P is minor-closed and positive on some $g \times g$ grid then it is positive on a graph G with treewidth $p > 20^{2g^5}$ via Theorem 1.3. Plugging this value of p into the bound from Theorem 4.3 gives the desired result. \square

Remark 4.2 *The results in [30, 23] can be used to obtain a singly exponential dependence on g as well, provided P can be shown to be positive on a graph that contains a grid-like minor of size g .*

4.2 Bounds for Erdos-Pósa theorems

Let \mathcal{F} be any family of graphs. Following the notation in [29], we say that the \mathcal{F} -packing number of G , denoted by $p_{\mathcal{F}}(G)$, is the maximum number of node-disjoint subgraphs of G , each of which is isomorphic to a member of \mathcal{F} . An \mathcal{F} -cover is a set X of vertices, such that $p_{\mathcal{F}}(G - X) = 0$; that is, removing X ensures that there is no subgraph isomorphic to a member of \mathcal{F} in G . The \mathcal{F} -covering number of G , denoted by $c_{\mathcal{F}}(G)$ is the minimum cardinality of an \mathcal{F} -cover for G . It is clear that $p_{\mathcal{F}}(G) \leq c_{\mathcal{F}}(G)$ always holds. A family \mathcal{F} is said to satisfy the Erdos-Pósa property if there is function $f : \mathbb{Z}_+ \rightarrow \mathbb{Z}_+$ such that $c_{\mathcal{F}}(G) \leq f(p_{\mathcal{F}}(G))$ for all graphs G . Erdos and Pósa [16] showed such a property when \mathcal{F} is the family of cycles, with $f(k) = \Theta(k \log k)$.

There is an important connection between treewidth and Erdos-Pósa property as captured by the following two lemmas. The proofs appear in the full version.

Lemma 4.3 *Let \mathcal{F} be any family of connected graphs, and let $h_{\mathcal{F}}$ be an integer-valued function, such that the following holds. For any integer k , and any graph G with $\text{tw}(G) \geq h_{\mathcal{F}}(k)$, G contains k disjoint subgraphs G_1, \dots, G_k , each of which is isomorphic to a member of \mathcal{F} . Then \mathcal{F} has the Erdos-Pósa property with $f_{\mathcal{F}}(k) \leq k \cdot h_{\mathcal{F}}(k)$.*

Lemma 4.4 *Let \mathcal{F} be any family of connected graphs, and let $h_{\mathcal{F}}$ be an integer-valued function, such that the following holds. For any integer k , and any graph G with $\text{tw}(G) \geq h_{\mathcal{F}}(k)$, G contains k disjoint subgraphs G_1, \dots, G_k , each of which is isomorphic to a member of \mathcal{F} . Moreover, suppose*

that $h_{\mathcal{F}}(\cdot)$ is superadditive⁶ and satisfies the property that $h_{\mathcal{F}}(k+1) \leq \alpha h_{\mathcal{F}}(k)$ for all $k \geq 1$ where α is some universal constant. Then \mathcal{F} has the Erdos-Pósa property with $f_{\mathcal{F}}(k) \leq \beta h_{\mathcal{F}}(k) \log(k+1)$ where β is a universal constant.

One way to prove that $p_{\mathcal{F}}(G) \geq k$ whenever $\text{tw}(G) \geq h_{\mathcal{F}}(k)$ is via the following proposition, that is based on the Grid-Minor Theorem. It is often implicitly used; see [29].

Proposition 4.1 *Let \mathcal{F} be any family of connected graphs, and assume that there is an integer g , such that any graph containing a $g \times g$ grid as a minor is guaranteed to contain a sub-graph isomorphic to a member of \mathcal{F} . Let $h(g')$ be the treewidth that guarantees the existence of a $g' \times g'$ grid minor in any graph. Then $f_{\mathcal{F}}(k) \leq O(k \cdot h(g\sqrt{k}))$. In particular $f_{\mathcal{F}}(k) \leq 2^{O(g^5 k^{2.5})}$.*

We improve the exponential dependence on k in the preceding proposition to near-linear. We state a more general theorem, whose proof appears in the full version, and then derive the improvement as a corollary.

Theorem 4.4 *Let \mathcal{F} be any family of connected graphs, and assume that there is an integer r , such that any graph of treewidth at least r is guaranteed to contain a sub-graph isomorphic to a member of \mathcal{F} . Then $f_{\mathcal{F}}(k) \leq \tilde{O}(kr^2)$.*

Corollary 4.2 *Let \mathcal{F} be any family of connected graphs, such that for some integer g , any graph containing a $g \times g$ grid as a minor is guaranteed to contain a sub-graph isomorphic to a member of \mathcal{F} . Then $f_{\mathcal{F}}(k) \leq 2^{O(g^5)} \tilde{O}(k)$.*

Some concrete results: For a fixed graph H , let $\mathcal{F}(H)$ be the family of all graphs that contain H as a minor. Robertson and Seymour [32], as one of the applications of their Grid-Minor Theorem, showed that $\mathcal{F}(H)$ has the Erdos-Pósa property iff H is planar. The if direction can be deduced as follows. Every planar graph H is a minor of a $g \times g$ grid, where $g = O(|V(H)|^2)$. We can then use Proposition 4.1 to obtain a bound on $f_{\mathcal{F}(H)}$, which is super-exponential in k . However, by directly applying Corollary 4.2, we get the following improved near-linear dependence on k .

Theorem 4.5 *For any fixed planar graph H , the family $\mathcal{F}(H)$ of graphs has the Erdos-Pósa property with $f_{\mathcal{F}(H)}(k) = O(k \cdot \text{poly} \log(k))$.*

For any integer $m > 0$, let \mathcal{F}_m be the family of all cycles whose length is 0 modulo m . Thomassen [35] showed that \mathcal{F}_m has the Erdos-Pósa property, with $f_{\mathcal{F}_m} = 2^{mO(k)}$. We can use Corollary 4.2 to obtain a bound of $f_{\mathcal{F}_m} = \tilde{O}(k) \cdot 2^{\text{poly}(m)}$, using the fact that a graph containing a grid minor of size $2^{\text{poly}(m)}$ must contain a cycle of length 0 modulo m .

Acknowledgements: CC thanks Tasos Sidiropoulos for discussions on grid minors and bidimensionality, and Alina Ene for discussions on the decomposition result in [8] and for pointing out [19]. This paper is partly inspired by the lovely article of Bruce Reed on treewidth and their applications [29]. We thank Paul Wollan for pointers to recent work on the Grid-Minor theorem and Erdos-Pósa type theorems.

⁶We say that an integer-valued function h is superadditive if for all $x, y \in \mathbb{Z}^+$, $h(x) + h(y) \leq h(x + y)$.

5. REFERENCES

- [1] E. Amir. Approximation algorithms for treewidth. *Algorithmica*, 56(4):448–479, 2010.
- [2] M. Andrews. Approximation algorithms for the edge-disjoint paths problem via Raecke decompositions. In *Proceedings of IEEE FOCS*, pages 277–286, 2010.
- [3] S. Arora, S. Rao, and U. V. Vazirani. Expander flows, geometric embeddings and graph partitioning. *J. ACM*, 56(2), 2009.
- [4] H. L. Bodlaender, R. Downey, F. V. Fomin, and D. Marx, editors. *The Multivariate Algorithmic Revolution and Beyond*. Essays Dedicated to Michael R. Fellows on the Occasion of His 60th Birthday. Springer-Verlag New York Incorporated, May 2012.
- [5] C. Chekuri and A. Ene. Poly-logarithmic approximation for maximum node disjoint paths with constant congestion. In *Proc. of ACM-SIAM SODA*, 2013.
- [6] C. Chekuri, S. Khanna, and F. B. Shepherd. The all-or-nothing multicommodity flow problem. In *Proc. of ACM STOC*, pages 156–165, 2004.
- [7] C. Chekuri, S. Khanna, and F. B. Shepherd. Multicommodity flow, well-linked terminals, and routing problems. In *Proc. of ACM STOC*, pages 183–192, 2005.
- [8] J. Chuzhoy. Routing in undirected graphs with constant congestion. In *Proc. of ACM STOC*, pages 855–874, 2012.
- [9] J. Chuzhoy and S. Li. A polylogarithmic approximation algorithm for edge-disjoint paths with congestion 2. In *Proc. of IEEE FOCS*, 2012.
- [10] E. Demaine and M. Hajiaghayi. The Bidimensionality Theory and Its Algorithmic Applications. *The Computer Journal*, 51(3):292–302, Nov. 2007.
- [11] E. Demaine and M. Hajiaghayi. Linearity of grid minors in treewidth with applications through bidimensionality. *Combinatorica*, 28:19–36, 2008.
- [12] E. Demaine, M. Hajiaghayi, and K.-i. Kawarabayashi. Algorithmic graph minor theory: Improved grid minor bounds and wagner’s contraction. *Algorithmica*, 54:142–180, 2009.
- [13] E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Subexponential parameterized algorithms on bounded-genus graphs and H-minor-free graphs. *J. ACM*, 52(6):866–893, 2005.
- [14] E. D. Demaine and M. Hajiaghayi. Quickly deciding minor-closed parameters in general graphs. *European Journal of Combinatorics*, 28(1):311–314, Jan. 2007.
- [15] R. Downey and M. Fellows. *Fixed-parameter Tractability*. Theory of computing systems. Springer, 2007.
- [16] P. Erdos and L. Pósa. On independent circuits contained in a graph. *Canadian Journal of Mathematics. Journal Canadien de Mathématiques*, 17:347–352, 1965.
- [17] U. Feige, M. Hajiaghayi, and J. Lee. Improved approximation algorithms for minimum weight vertex separators. *SIAM J. Comput.*, 38:629–657, 2008.
- [18] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. an Eats Series. Springer, 2010.
- [19] F. V. Fomin, S. Saurabh, and D. M. Thilikos. Strengthening Erdos-Pósa property for minor-closed graph classes. *Journal of Graph Theory*, 66(3):235–240, 2011.
- [20] M. Grohe and D. Marx. On tree width, bramble size, and expansion. *J. Comb. Theory, Ser. B*, 99(1):218–228, 2009.
- [21] K. Kawarabayashi and Y. Kobayashi. Linear min-max relation between the treewidth of H-minor-free graphs and its largest grid minor. In *Proc. of STACS*, 2012.
- [22] R. Khandekar, S. Rao, and U. Vazirani. Graph partitioning using single commodity flows. *J. ACM*, 56(4):19:1–19:15, July 2009.
- [23] S. Kreutzer and S. Tazari. On brambles, grid-like minors, and parameterized intractability of monadic second-order logic. In *Proc. of ACM-SIAM SODA*, pages 354–364, 2010.
- [24] F. T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46:787–832, 1999.
- [25] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15:215–245, 1995.
- [26] R. Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford Lecture Series in Mathematics And Its Applications. Oxford University Press, 2006.
- [27] H. Räcke. Minimizing congestion in general networks. In *Proc. of IEEE FOCS*, pages 43–52, 2002.
- [28] S. Rao and S. Zhou. Edge disjoint paths in moderately connected graphs. *SIAM J. Comput.*, 39(5):1856–1887, 2010.
- [29] B. Reed. *Surveys in Combinatorics*, chapter Treewidth and Tangles: A New Connectivity Measure and Some Applications. London Mathematical Society Lecture Note Series. Cambridge University Press, 1997.
- [30] B. A. Reed and D. R. Wood. Polynomial treewidth forces a large grid-like-minor. *European Journal of Combinatorics*, 33(3):374–379, Apr. 2012.
- [31] N. Robertson, P. Seymour, and R. Thomas. Quickly Excluding a Planar Graph. *Journal of Combinatorial Theory, Series B*, 62(2):323–348, Nov. 1994.
- [32] N. Robertson and P. D. Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, Aug. 1986.
- [33] P. Seymour. The tree-width bound from excluding a planar minor. Third Workshop on Graphs and Matroids, August 2012.
- [34] P. D. Seymour and R. Thomas. Graph searching and a min-max theorem for tree-width. *J. Comb. Theory Ser. B*, 58(1):22–33, May 1993.
- [35] C. Thomassen. On the presence of disjoint subgraphs of a specified type. *Journal of Graph Theory*, 12(1):101–111, 1988.