

An Improvement of Reed’s Treewidth Approximation

Mahdi Belbasi Martin Fürer

Computer Science and Engineering Department, Pennsylvania State University
University Park, PA USA

Submitted: July 2021	Reviewed: November 2021	Revised: March 2022
Reviewed: May 2022	Revised: June 2022	Accepted: June 2022
Final: June 2022	Published: June 2022	
Article type: Regular Paper		
Communicated by: S.-H. Hong, S. C. Nandy, R. Uehara		

Abstract. We present a new approximation algorithm for the treewidth problem which finds an upper bound on the treewidth and constructs a corresponding tree decomposition as well. Our algorithm is a faster variation of Reed’s classical algorithm. For the benefit of the reader, and to be able to compare these two algorithms, we start with a detailed time analysis of Reed’s algorithm. We fill in many details that have been omitted in Reed’s paper. Computing tree decompositions parameterized by the treewidth k is fixed parameter tractable (FPT), meaning that there are algorithms running in time $\mathcal{O}(f(k)g(n))$ where f is a computable function, and $g(n)$ is polynomial in n , where n is the number of vertices. An analysis of Reed’s algorithm shows $f(k) = 2^{\mathcal{O}(k \log k)}$ and $g(n) = n \log n$ for a 5-approximation. Reed simply claims time $\mathcal{O}(n \log n)$ for bounded k for his constant factor approximation algorithm, but the bound of $2^{\Omega(k \log k)} n \log n$ is well known. From a practical point of view, we notice that the time of Reed’s algorithm also contains a term of $\mathcal{O}(k^2 2^{24k} n \log n)$, which for small k is much worse than the asymptotically leading term of $2^{\mathcal{O}(k \log k)} n \log n$. We analyze $f(k)$ more precisely, because the purpose of this paper is to improve the running times for all reasonably small values of k .

Our algorithm runs in $\mathcal{O}(f(k)n \log n)$ too, but with a much smaller dependence on k . In our case, $f(k) = 2^{\mathcal{O}(k)}$. This algorithm is simple and fast, especially for small values of k . We should mention that Bodlaender et al. [2016] have an algorithm with a linear dependence on n , and Korhonen [2021] obtains the much better approximation ratio of 2, while the current paper achieves a better dependence on k .

Special Issue on the 15th Int. Conference and Workshops on Algorithms and Computation, WALCOM 2021

E-mail addresses: belbasi@psu.edu (Mahdi Belbasi) fhs@psu.edu (Martin Fürer)



This work is licensed under the terms of the [CC-BY](https://creativecommons.org/licenses/by/4.0/) license.

1 Introduction

Since the 1970s and early 1980s, when the notions of treewidth and tree decomposition were introduced [5, 13, 19], they have played important roles in computer science [7]. In a nutshell, treewidth is a parameter of a graph that measures how similar it is to a tree. One of the main reasons that the tree decomposition is widely studied is that many NP-complete problems have efficient algorithms for graphs with small treewidth. In fact, Courcelle's metatheorem [9] states that every graph property definable in monadic second-order logic of graphs can be solved in linear time on graphs of bounded treewidth. The first step of solving such problems is to find an optimal or near-optimal tree decomposition. However, finding an optimal tree decomposition itself is NP-hard [2]. In this work, we propose an algorithm which is based on Reed's algorithm [17] to approximate the treewidth and find an approximately optimal tree decomposition.

Definition 1 *A graph problem is fixed parameter tractable (FPT) if it can be solved in time $\mathcal{O}(f(k)n^{\mathcal{O}(1)})$, where f is a computable function, k is a parameter of the graph, and n is the input graph size.*

1.1 Previously Known Results

In this work, we are interested in algorithms which run fast (polynomial in terms of the number of vertices) for graphs with bounded treewidth. One of the first algorithms given for this problem goes back to the same paper where treewidth has been shown to be NP-complete. Arnborg et al. [2] gave an algorithm which runs in time $\mathcal{O}(n^{k+2})$. In 1995, Robertson and Seymour [20] gave a quadratic time FPT approximation algorithm. Later on, Lagergren introduced an 8-approximation algorithm with the time complexity of $2^{\mathcal{O}(k \log k)} n \log^2 n$ [16]. In 1992, Reed [17] improved these algorithms to have an algorithm running in time $2^{\mathcal{O}(k \log k)} n \log n$. In this paper, we formally show that the approximation ratio of Reed's algorithm is 7 or 5, depending on the frequency of the split by volume. We show that this algorithm runs in time $\mathcal{O}(2^{2^{4k}} k! n \log n)$, in order to be able to compare it to our algorithm. Like the algorithm of Robertson and Seymour, Reed's algorithm is recursive. In [20], they find a separator that partitions G into two parts but they do not force the separator to partition the entire graph in a balanced fashion. Reed finds a separator which partitions the graph in a balanced way to obtain time $\mathcal{O}(n \log n)$ for bounded k . This paper focuses on this algorithm. Later, Bodlaender gave an exact algorithm which runs in $2^{\mathcal{O}(k^3)} n$ [6]. Even though we focus only on constant-factor approximation algorithms, it is worth mentioning the $(\log k)$ -approximation algorithm by Amir [1] and the $\sqrt{\log k}$ -approximation algorithm by Feige et al. [10]. Later in 2016, Bodlaender et al. [8] gave two constant factor approximation algorithms which run in $2^{\mathcal{O}(k)} \mathcal{O}(n \log n)$ and $2^{\mathcal{O}(k)} \mathcal{O}(n)$ respectively. The former is a 3-approximation and the latter is a 5-approximation. Although it is a great result from a theoretical point of view, it uses a sophisticated data structure and the constant factor hidden in $\mathcal{O}(k)$ in the exponent is not claimed to be practical. Very recently, Korhonen gave a 2-approximation algorithm for the same problem running in time $2^{\mathcal{O}(k)} n$ [15]. He first provides a loose upper bound on the hidden coefficient of k in the exponent. Then he improves it to 10.7549 by decreasing his potential function to what seems to be a natural barrier in the worst-case scenario. Here, we sacrifice the linear dependence on n , but drop the coefficient of k in the exponent to only 7.61. Table 1.1 summarizes the history of previously known algorithms for treewidth problem.

Reference	Approximation Ratio	Dependence on k	Dependence on n	Comments
Arnborg et al. (1987) [3]	1	$\mathcal{O}(n^{k+2})$		not FPT
Robertson & Seymour (1995) [18, 20]	4	$\mathcal{O}(3^{3k})$	n^2	
Lagergren (1996) [16]	8	$2^{\mathcal{O}(k \log k)}$	$n \log^2 n$	
Reed (1992) [17]	5 (or 7)	$2^{24k} k!$ (Sec. 3)	$\mathcal{O}(n \log n)$	
Bodlaender (1996) [6]	1	$2^{\mathcal{O}(k^3)}$	n	
Feige et al. (2008)	$\mathcal{O}(\sqrt{\log k})$	$\mathcal{O}(1)$	$n^{\mathcal{O}(1)}$	not a constant-factor approximation
Amir (2010) [1]	$\mathcal{O}(\log k)$	$\mathcal{O}(k \log k)$	n^4	not a constant-factor approximation
	4	$\mathcal{O}(2^{4.38k} k)$	n^2	
	4.5	$\mathcal{O}(2^{3k} k^{1.5})$	n^2	
	$\frac{11}{3}$	$\mathcal{O}(2^{3.6982k} k^3)$	$n^3 \log^4 n$	
Fomin et al. (2015) [12]	1	$\mathcal{O}(1)$	1.7347^n	not FPT
Bodlaender et al. (2016) [8]	3	$2^{\mathcal{O}(k)}$	$n \log n$	the coefficients of k are not stated
	5	$2^{\mathcal{O}(k)}$	n	
Belbasi & Fürer (2021) [4]	5	$2^{8.766k}$	$n \log n$	
Korhonen (2021) [15]	2	$2^{10.7549k}$	n	relatively low coefficient of k in the exponent, good approximation ratio, and only linear
This paper	5	$2^{7.61k}$	$n \log n$	extra $\log n$ in the running time compared to [15] but has smaller coefficient of k in the exponent

Table 1: The history of previous algorithms for the treewidth approximation problem

1.2 Our Contribution

First in Section 3, we analyze Reed's algorithm [17] in detail. Reed has focused on the dependence on n because his goal was to come up with an $\mathcal{O}(n \log n)$ -time algorithm, for fixed k . The "fastest"¹ algorithm at that time, was Lagergren's $\mathcal{O}(n \log^2 n)$ -time algorithm [16]. We show that the dependence on k in Reed's algorithm is $2^{\mathcal{O}(k \log k)}$. Furthermore, we give a proof for the approximation ratio of Reed's algorithm by filling in the details.

Then, we propose two improvements and prove that the approximation ratio stays at 5. One of our improvements focuses on the notion of a "balanced split". We call a split balanced, if we get two parts of volume $1 - \epsilon$ and ϵ (or better). Then, the running time of our algorithm has a factor of $1/\epsilon$. For instance, if we set $\epsilon = \frac{1}{100}$, a generous estimation shows that the dependence on k in our $\mathcal{O}(f(k)n \log n)$ -time algorithm is $k^2 2^{8.87k}$, instead of $2^{24k}(k+1)!$ in Reed's algorithm. Here the asymptotic notation is a bit misleading from a practical point of view, as $2^{24k} = o(k!)$, even though $k!$ is reasonable for small k , while 2^{24k} is not. Then, we further improve the running time upper bound and show that our algorithm runs in time $\mathcal{O}(2^{7.61k} k^2 n \log n)$. In the end, the main aim of this paper is to produce a constant-factor approximation algorithm that runs in time $2^{ck} n \log n$ with c as small as possible.

2 Preliminaries

2.1 Tree Decomposition

Definition 2 A tree decomposition of a graph $G = (V, E)$, is a tree $\mathcal{T} = (V_{\mathcal{T}}, E_{\mathcal{T}})$ such that each node x in $V_{\mathcal{T}}$ is associated with a set B_x (called the bag of x) of vertices in G , and such that \mathcal{T} has the following properties:

- The union of all bags is equal to V . In other words, for each $v \in V$, there exists at least one node $x \in V_{\mathcal{T}}$ with B_x containing v .
- For every edge $\{u, v\} \in E$, there exists a node x such that $u, v \in B_x$.
- For any nodes $x, y \in V_{\mathcal{T}}$, and any node $z \in V_{\mathcal{T}}$ belonging to the path connecting x and y in \mathcal{T} , $B_x \cap B_y \subseteq B_z$.

In this paper we use a variation of tree decomposition where the adjacent bags differ in at most one vertex (converting can happen in linear time).

The *width of a tree decomposition* is the size of a largest bag minus one. The *treewidth* of a graph G is the minimum width over all tree decompositions of G called $tw(G)$. Observe that the treewidth of a tree is 1. In the following, we reserve the letter k for the treewidth+1.

We have to mention that Bodlaender et al. [8] filled in some details on Reed's algorithm. We need to be more detailed because we do not use Reed's algorithm as a black box. That is why we first analyze Reed's algorithm precisely (Section 3), fill in the blanks, and then introduce our improvements of his algorithm (Section 4).

¹Considering the dependence on both n and k .

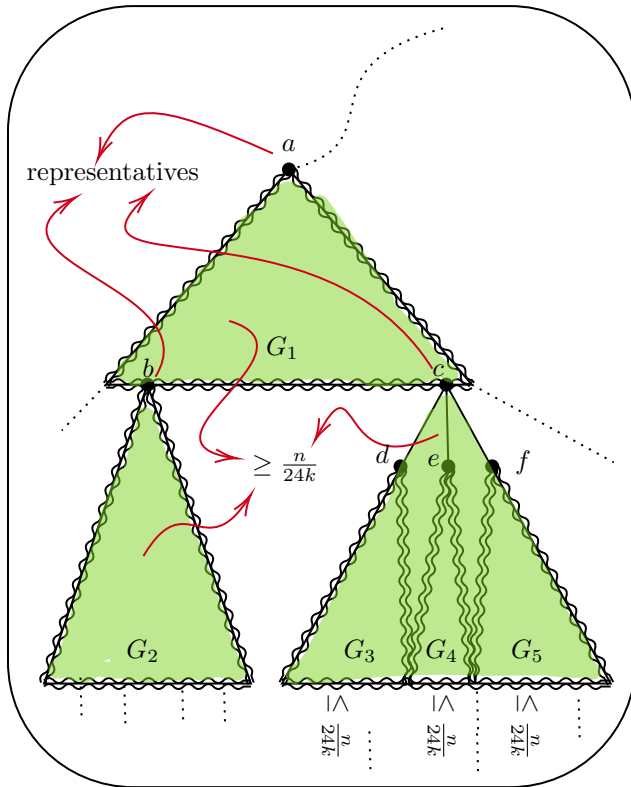


Figure 1: For a given graph G , let $V(G)$ be the set of its vertices. Let $G_6 = G[V(G_3) \cup V(G_4) \cup V(G_5) \cup \{c\}]$ (the entire rightmost green subtree rooted at c), and assume $|V(G_1)|, |V(G_2)|$, and $|V(G_6)| \geq \frac{n}{24k}$, and $|V(G_3)|, |V(G_4)|$, and $|V(G_5)| < \frac{n}{24k}$. Here, d, e , and f are NOT representatives but a, b , and c are.

3 Analysis of Reed’s Algorithm

In 1992, Reed gave an elegant algorithm [17] to either construct a tree decomposition of width less than $7k$ or $5k$ of a given graph G , or declare that the treewidth is at least k and output a subgraph which is a bottleneck (no separator of size $\leq k$).

3.1 Summary of Reed’s Algorithm

In Reed’s algorithm, one of the main tasks is to find a “balanced” separator S that splits the graph $G - S$ into two subgraphs with sets of vertices $X, Y \subseteq V(G)$. Once a balanced separator is found, the algorithm recursively finds a tree decomposition for $G[X \cup S]$ (the subgraph induced by $X \cup S$) and $G[Y \cup S]$.

The main task is to find a balanced separator. Instead of branching on every vertex (going to X, Y , or S , which would be exponential in n), Reed forms groups of vertices and works with the representatives of the groups. Then, he branches on the representatives.




Figure 2: Error, i.e. part of a subtree not on the side of its representative

Reed does a DFS and finds the deepest vertex v whose subtree has at least $\frac{n}{24k}$ vertices² (See Figure 1). We call such a vertex a “representative”. He defines the weight of the representative v as the size of its subtree rooted at v in the DFS, denoted by $w(v)$. The idea here, is that if some representatives with a large total weight go to either X or Y , then most of their descendants will go to the same set. The reason is that if a descendant goes to the other side, the path connecting the representative to the descendant has at least one vertex in the separator. However, we know that the separator cannot have more than k vertices. So, not many vertices will go to the wrong set; not more than $\frac{n}{24}$ vertices, in total. This is because every subtree that partially goes to the other side should go through the separator and have one vertex there. Hence, not more than k subtrees rooted at children of the representatives can go through the separator, which results in at most $\frac{n}{24}$ vertices on the wrong side. This nice property allows Reed to work with the set of representatives (which is much smaller) rather than all the vertices (see Figure 2).

Now, one might think, why not just check all the possibilities of the representatives going to X, Y , or S ? The reason that this simple idea does not work is that if a representative goes to the separator, its entire subtree of arbitrary size can go anywhere and we do not have any control over it. Reed handles this problem by deciding whether any representative is going to the separator, at the very beginning of the algorithm. If so, he just places such a representative (namely v) into S (and not its subtree) and starts forming a new group of representatives by running a new DFS on $G - \{v\}$. So, the other representatives might change. Also, since one vertex has been placed into the separator, now $k \leftarrow k - 1$. However, if none of the representatives goes to the separator, he

²Later, we talk about this threshold.

branches on placing them left (X) or right (Y). This is the high-level idea of Reed’s algorithm. Let’s start with presenting and reviewing some definitions.

3.2 Centroids and Separators

For an undirected graph $G = (V, E)$ and a subset W of the vertices, $G[W]$ is the subgraph induced by W . For the sake of simplicity throughout this paper, let $G - W$ be $G[V \setminus W]$ and $G - v$ be $G - \{v\}$ for any $W \subseteq V(G)$ and any $v \in V(G)$.

Also, in a weighted graph, a non-negative integer weight $w(v)$ is defined for each vertex v . For a subset W of the vertices, the weight $w(W)$ is simply the sum of the weights of all vertices in W . Furthermore, the total weight or the weight of G is the weight of V .

Definition 3 *A centroid of a weighted tree T is a node x such that none of the trees in the forest $T - x$ has more than half the total weight.*

Definition 4 *A tree decomposition is called a good tree decomposition if the adjacent bags differ in at most one vertex.*

For good tree decompositions we choose a stronger version of centroid.

Lemma 1 *Any tree decomposition \mathcal{T} with width k can be converted in linear time to a good tree decomposition \mathcal{T}^* of the same graph without increasing the width.*

Proof: Bodlaender [6] has defined the more restrictive notion of a smooth tree decomposition and shown that it can be computed in linear time. However, we include the derivation for our notion to make it self-contained³.

Let x and y be two adjacent nodes of \mathcal{T} with associated bags B_x , and B_y , respectively.

Let $B_x = \{v_1, \dots, v_t\}$, and $B_y = \{u_1, \dots, u_s\}$. Notice that $s, t \leq k$.

Now, let $V^- = B_x \setminus B_y$, and $U^+ = B_y \setminus B_x$. Let $D = V^- \cup U^+ = \{d_1, \dots, d_{|V^-|}, d'_1, \dots, d'_{|U^+|}\}$.

We just add $|D|$ intermediate nodes (namely $x_1, x_2, \dots, x_{|V^-|}, y_1, y_2, \dots, y_{|U^+|}$) between x and y such that we start by deleting vertices of V^- one by one and then adding vertices of U^+ one at a time.

Set

$$\begin{cases} B_{x_1} = B_x \setminus \{d_1\} \\ B_{x_i} = B_{x_{i-1}} \setminus \{d_i\} & \forall i \in \{2, \dots, |V^-|\} \\ B_{y_1} = B_{x_{|V^-|}} \cup \{d'_1\} \\ B_{y_i} = B_{y_{i-1}} \cup \{d'_i\} & \forall i \in \{2, \dots, |U^+|\}. \end{cases} \tag{1}$$

□

□

Definition 5 *A strong centroid of a good tree decomposition τ of a graph $G = (V, E)$ with respect to $W \subseteq V$ is a node x of τ such that none of the connected components of $G - B_x$ contains more than $\frac{1}{2}|W \setminus B_x|$ vertices of W .*

The following lemma shows the existence of a strong centroid for any given subset W of V .

³Alternatively, one can refer to the derivation of a similar result for nice tree decomposition [14].

Lemma 2 *For every good tree decomposition $(\mathcal{T}, \{B_x : x \in V_{\mathcal{T}}\})$ of a graph $G = (V, E)$ and every subset $W \subseteq V$, there exist a strong centroid with respect to W ⁴.*

Proof: If a node x is not a strong centroid with respect to W , then let C_x be the set of vertices in the unique connected component of $G - B_x$ containing more than $\frac{1}{2}|W \setminus B_x|$ vertices of X . In the forest obtained from the tree \mathcal{T} by removing x , there is a tree \mathcal{T}_x with the property that the union of all bags in \mathcal{T}_x contains all the vertices of C_x .

Now, we define a set F of directed tree edges by $(x, y) \in F$ if all the following conditions hold:

- x is not a strong centroid
- y is a neighbor of x in \mathcal{T} ,
- y is a node in \mathcal{T}_x .

Now we show that there is a node x with out-degree 0 in $(V_{\mathcal{T}}, F)$. Such an x is a centroid, and we are done. Otherwise, F contains (x, y) and (y, x) for some $x, y \in V_{\mathcal{T}}$. W.l.o.g., $B_y = B_x \cup \{v\}$ for some $v \in V \setminus B_x$. Note that \mathcal{T}_x and \mathcal{T}_y are disjoint. Furthermore, any vertex that is in a bag of \mathcal{T}_x and in a bag of \mathcal{T}_y is also in B_x and B_y . Thus also C_x and C_y are disjoint.

Furthermore, $W \cap C_x \subseteq W \setminus B_x$ and $W \cap C_y \subseteq W \setminus B_y \subseteq W \setminus B_x$. As $W \cap C_x$ and $W \cap C_y$ are disjoint, they cannot both have more than $\frac{1}{2}|W \setminus B_x|$ vertices, which is a contradiction.

Hence, there exists a node x which is a strong centroid. □

□

We use the definitions of balanced W -separator and weakly balanced W separation from the book of Flum and Grohe [11].

Definition 6 *Let $G = (V, E)$ be a graph and $W \subseteq V$. A balanced W -separator is a set $S \subseteq V$ such that every connected component of $G - S$ has at most $\frac{1}{2}|W|$ vertices.*

Lemma 3 [11, Lemma 11.16] *Let $G = (V, E)$ be a graph of treewidth at most $k - 1$ and $W \subseteq V$. Then there exists a balanced W -separator of G of size at most k .*

We say that a separator S separates $X \subseteq V$ from $Y \subseteq V$ if $C \cap X = \emptyset$ or $C \cap Y = \emptyset$ for every connected component C of $G - S$.

Definition 7 *Let $G = (V, E)$ be a graph and $W \subseteq V$. A weakly balanced separation of W is a triple (X, S, Y) , where $X, Y \subseteq W$, $S \subseteq V$ are pairwise disjoint sets such that:*

- $|S| \leq k$
- $W = X \cup (S \cap W) \cup Y$
- S separates X from Y
- $0 < |X|, |Y| \leq \frac{2}{3}|W|$.

⁴Flum-Grohe [11] use the property of a standard centroid for tree decompositions in Lemma 11.16 page 267. They talk about balanced W -separators and do not use the term centroid. They only use the standard notion of a centroid, showing that no connected component contains more than $\frac{|W|}{2}$ of the vertices. Strong centroids show that no connected component has size more than $\frac{|W \setminus S|}{2}$.

Lemma 4 [11, Lemma 11.19] *For $k \geq 3$, let $G = (V, E)$ be a graph of treewidth at most $k - 1$ and $W \subseteq V$ with $|W| \geq 2k + 1$. Then there exists a weakly balanced separation of W of size at most k .*

Even though Lemma 4 is sufficient for us, one can make it stronger such that it holds for $|W| \geq k + 1$.

Theorem 1 [11, Corollary 11.22] *For a graph of treewidth at most $k - 1$ with a given set $W \subseteq V$ of size $|W| = 3k - 2$, a weakly balanced separation of W of size $\leq k$ can be found in time $O(3^{3k} k^2 n)$.*

3.3 Algorithm to Find a Weakly Balanced Separation

Separation(G, k) is the main part of Reed’s algorithm. It finds a separator of size at most k in G using the procedures Split(G, X, Y, k) and DFS-Trees(G, k). We explain each of these procedures.

3.3.1 Split(G, X, Y, k)

For X, Y disjoint subsets of V , Split(G, X, Y, k) finds a separator S of size at most k in G which is strictly between X and Y . Split reports failure if no such separator exists (described in Lemma 11.20 of [11]).

3.3.2 DFS-Trees(G, k)

DFS-Trees(G, k) (Algorithm 1 using Algorithm 2) computes a DFS tree and partitions it into smaller DFS trees with the following properties.

- the size (number of vertices) of the smaller trees is at least $s = n/24k$, and
- all subtrees rooted at children of the roots of the trees in the partition have size less than $n/24k$.

DFS-Trees(G, k) collects the set W' consisting of all the roots of the trees in the partition. These roots are representatives of the vertices in their small DFS tree. Therefore, the weight $w[v]$ for $v \in W'$ is the number of vertices in the small tree with root v .

3.3.3 Separation(G, k)

Separation(G, k) is the recursive procedure that splits according to the number of vertices (Algorithm 3). Note that when any vertex v is placed into the separator S , then the procedure Separation removes that vertex v from the graph and starts from scratch. The idea is that when we place a root of a small tree (a representative) left or right, then we want to put the whole small tree there. But when a representative is placed into the separator, then its tree does not go there. At this point a new collection of trees is formed.

3.4 The Correctness of Reed’s Algorithm

If the treewidth is at most $k - 1$, then there is a good tree decomposition of G of width $k - 1$. Let x be a centroid in it. The connected components of $G[V \setminus B_x]$ can be partitioned into 2 parts L and R , such that no part has more than $\frac{2}{3}|V|$ vertices.

First, we prove the correctness of the algorithm for the case that G is a connected graph (lines 16-25 of Algorithm 3). Later, we describe the case that G is not connected.

Algorithm 1: Construct Small DFS-Trees

Result: Roots of DFS-Trees (representatives) with sizes of their strict subtrees
 $< |V|/(24k)$

```

1 Procedure DFS-Trees( $G, k$ )                                //  $G$  is a connected graph.
2 *  $s = \frac{|V|}{24k}$                                            //  $s$ : the size bound for splitting off a small tree.
3 *  $W' = \emptyset$ 
4 * for all  $v \in V$  do
5   | *  $\text{color}[v] = \text{WHITE}$ 
6 end
7 * Pick any vertex  $u$  of  $G$ .
8 *  $\text{count} = \text{DFS-visit}(G, u)$ 
9 *  $w[x] = w[x] + \text{count}$ , where  $x$  is the last representative added to  $W'$ 
10 * return ( $W'$ ,  $w[v]$  for all  $v \in W'$ )
11 End Procedure

```

Algorithm 2: Main recursive procedure of DFS-Trees

```

1 Procedure DFS-Visit( $G, u$ )
2 *  $\text{color}[u] = \text{GRAY}$ 
3 *  $\text{count} = 1$ 
4 * for all  $v$  adjacent to  $u$  do
5   | * if  $\text{color}[v] == \text{WHITE}$                                 // The white vertex  $v$  is discovered now.
6   |   then
7   |     | *  $\text{count} = \text{count} + \text{DFS-Visit}(G, v)$ 
8   |     end
9   end
10 * if  $\text{count} \geq s$  then
11   | *  $W' = W' \cup \{u\}$ 
12   | *  $w[u] = \text{count}$ 
13   | *  $\text{count} = 0$ 
14 end
15 return  $\text{count}$ 
16 End Procedure

```

Algorithm 3: Main recursive procedure in Reed’s algorithm

Result: A weakly balanced separation (X, S, Y) of G of size $\leq k$

```

1 Procedure SEPARATION( $G, k$ )
2 *if  $G$  is not a connected graph then
3   | *Let  $C_1, \dots, C_t$  be the connected components of  $G$ , and w.l.o.g., assume  $C_1$  is the
   | largest component.
4   | *if  $|C_1| < \frac{3}{4}|V|$  then
5   |   | *Let  $L = C_1$  and  $i = 2$ .
6   |   | *while  $|L| < \frac{1}{4}|V|$  do
7   |   |   | *L =  $L \cup C_i$ .
8   |   |   | *i =  $i + 1$ 
9   |   | end
10  |   | *return  $(L, \emptyset, V \setminus L)$ 
11  | end
12  | *Let  $(X', S', Y') = \text{SEPARATION}(G[C_1], k)$ . W.l.o.g., assume  $X'$  is the one with lower
   | weight.
13  | *return  $(X' \cup (\bigcup_{i=2}^t C_i), S', Y')$ 
14 end
15 *if  $k > 0$  then
16 | *(W', w[v] : \forall v \in W') = \text{DFS-Trees}(G, k)
17 end
18 *for all  $v \in W'$  // Here  $v$  is placed into separator  $S$ .
19 do
20 | *(X, S, Y) = \text{SEPARATION}(G - v, k - 1)
21 | *if  $\neg \text{failure}$  then
22 | | *return  $(X, S \cup \{v\}, Y)$ 
23 | end
24 end
   | // The set of vertices  $W'$  is partitioned into  $X \subseteq L$  and  $Y \subseteq R = W' \setminus L$ .
25 *for all  $X \subseteq W'$  // Here no vertex is put into  $S$ .
26 do
27 | *if  $(\frac{1}{3} - \frac{1}{24})|V| \leq w(X) \leq (\frac{2}{3} + \frac{1}{24})|V|$  then
28 | | *Split $(G, X, W' \setminus X, k)$ 
29 | | *if  $\neg \text{failure}$  then
30 | | | *return  $(X, S, Y)$ 
31 | | end
32 | end
33 end
34 return failure
35 End Procedure

```

Note that for the correctness proof, we do not have to find this tree decomposition. It is sufficient to know that it exists. We can assume, that we have fixed such a tree decomposition, a centroid x and the sets L and R .

Every set $W' \subseteq V$ is partitioned into parts in L , the separator $S = B_x$, and R .

One of the many branches of the procedure $\text{Separation}(G, k)$ working with a set W' , tries this partition of it, and succeeds, unless a previously taken branch has already succeeded. First, the procedure decides which part of W' goes into S , one vertex v at a time. This vertex v is removed from G , but otherwise, we still consider the same tree decomposition. $|B_x|$ has now decreased by 1, as v is removed from it.

We then consider the case that none of the remaining vertices in W' are in the separator. The procedure decides which part $X \subseteq W'$ goes into L (line 26) of Algorithm 3. Now the weight of each part is at most $(\frac{2}{3} + \frac{1}{24})n$ as at most k small subtrees can have some of their vertices on the wrong side. And this is at most k times less than $\frac{n}{24k}$ vertices.

On the branch of the procedure $\text{Separation}(G, k)$ which selects this partition of W' , there is the separator B_x of size at most k between X and Y . Our algorithm cannot guarantee to find this separator B_x , but it will find some separator S of size at most k between X and Y . Again up to $\frac{1}{24}n$ vertices can be on the opposite side of their representatives. Now the larger side can contain at most $(\frac{2}{3} + 2 \cdot \frac{1}{24})n = \frac{3}{4}n$ vertices. Thus we have a somewhat balanced partition (a constant fraction on each side).

Note that if at any point G becomes disconnected in Algorithm 3, it only works to our advantage, and we handle it separately in lines 2-15. Let C_1, \dots, C_t be the connected components of G , and w.l.o.g., assume C_1 has the highest volume. In this case, there are two possibilities:

- First, we consider the case that C_1 has volume (actual size) $\leq \frac{3}{4}n$. Then, if $|C_1|$ is already $\geq \frac{1}{4}n$ we are done. We return C_1 as the L.H.S. and the remaining components as the R.H.S. However, if $|C_1| < \frac{1}{4}n$, we add other components to C_1 until its volume passes $\frac{1}{4}n$ for the first time. Note that since $|C_1|$ has maximum volume, all other components have volume $\leq \frac{1}{4}n$, and adding them one at a time will not cause an issue.
- In the second case, C_1 has volume more than $\frac{3}{4}n$. Then, we run **SEPARATION** $(G[C_1], k)$ to get (X', S', Y') . This is a separation for $G[C_1]$. W.l.o.g., assume X' does not have a higher weight than Y' . Then, we put all the vertices represented by X' along with all other components $(X' \cup (\bigcup_{i=2}^t C_i))$ on the L.H.S. Note that:

$$\text{Vol}(X' \cup (\bigcup_{i=2}^t C_i)) = \text{Vol}(X') + n - |C_1|.$$

We know that $(\frac{1}{3} - \frac{1}{24})|C_1| \leq \text{Vol}(X') \leq (\frac{1}{2} + \frac{1}{24})|C_1|$. Hence,

$$\begin{aligned} \text{Vol}(X' \cup (\bigcup_{i=2}^t C_i)) &\leq (\frac{1}{2} + \frac{1}{24})|C_1| + n - |C_1| = n - \frac{13}{24}|C_1| \\ &\leq \frac{19}{32}n < \frac{3}{4}n. \end{aligned}$$

and

$$\text{Vol}(X' \cup (\bigcup_{i=2}^t C_i)) \geq (\frac{1}{3} - \frac{1}{24})|C_1| + n - |C_1| = n - \frac{17}{24}|C_1|$$

$$\geq n - \frac{17}{24}n = \frac{7}{24}n > \frac{1}{4}n.$$

Furthermore,

$$\text{Vol}(Y') \leq \frac{3}{4}|C_1| \leq \frac{3}{4}n,$$

and

$$\text{Vol}(Y') \geq \left(\frac{1}{2} - \frac{1}{12}\right)|C_1| \geq \frac{5}{12} \cdot \frac{3}{4}n = \frac{5}{16}n > \frac{1}{4}n.$$

So, in the case that G is not connected, we still find a somewhat balanced separator.

We assume the $\mathcal{O}(3^{3k}k^2n^2)$ 4-approximation algorithm of Robertson-Seymour [18, 20] (see Proposition 11.14 of Flum and Grohe [11]) is known. It handles a set $W \subseteq V$ of size $3k - 2$. Working with (G, W) , it finds a separator S of size at most k to split $G - S$ into two parts L and R with both, $|W \cap X|$ and $|W \cap Y|$ at most $\frac{2}{3}|W|$. A tree node with bag $W \cup S$ is formed. Two recursive calls continue with $(G[L \cup S], (W \cap L) \cup S)$ and $(G[R \cup S], (W \cap R) \cup S)$, respectively.

Reed’s algorithm does the same steps to handle W . However, in order to decrease the dependence on n of the running time from $\mathcal{O}(n^2)$ to $\mathcal{O}(n \log n)$, Reed intersperses these balanced partitions of W with balanced partitions of V (Algorithm 3). In each case, it would be desirable that V and W simultaneously split in a balanced way. However, during the traditional splitting of W , the graph might be split very unbalanced, and during the new splitting of V , the set W might not be split at all.

Reed’s algorithm can alternate between splitting W as in the $\mathcal{O}(3^{3k}k^2n^2)$ algorithm and splitting V . Now W can be of size at most $6k$. On each side, we have at most $(\frac{2}{3} \cdot 6k) = 4k$. Splitting by W as well as splitting by V adds k to the new W . Thus, we are back at $6k$. The constructed tree decomposition has then width at most $7k$. But we show that this can be improved to a 5-approximation algorithm. We do not need to alternate between splitting W and V . Splitting V is a costly procedure. We can do it only after every $\log_{\frac{3}{2}} k$ steps and we still spend only $\mathcal{O}(f(k)n \log n)$ time.

We start with W of size at most $4k$ ($3k$ and k_{excess} as excess). Initially, $k_{\text{excess}} = k$. Each time we split W , we get

$$|W| \leq \frac{2}{3} \cdot 3k + \underbrace{k}_{\text{adding separator}} + \frac{2}{3}k_{\text{excess}} = 3k + \frac{2}{3}k_{\text{excess}},$$

and then update $k_{\text{excess}} \leftarrow \lfloor \frac{2}{3}k_{\text{excess}} \rfloor$. The excess drops by a factor of $\frac{2}{3}$. After $\log_{\frac{3}{2}} k$ steps, the excess becomes zero and then we can split by V . At this point, $|W|$ could increase to $4k$ again ($3k$ was the size of W before this step, and when we split by V , we have to include the separator as well). We end up with $|W| \leq 4k$ and we add the separator to the root bag, which means the largest bag has size at most $5k$. Therefore, it is a 5-approximation algorithm. Reed mentions $5k$ in his paper, but he does not mention the frequency of the two operations. We don’t know whether he had the same modification in mind. Simple alternation between the two operations only achieves $7k$.

3.5 Running Time of Reed’s Algorithm

Let $T(n, k)$ be the running time of the procedure $SEPARATION(G, k)$ for $G = (V, E)$ and $n = |V|$. Let n' and k' be the current bound on the graph size and current separator capacity.

Initially $n' = n$ and $k' = k$. We have the following recurrence for Reed's algorithm.

$$T(n', k') \leq 24k'T(n' - 1, k' - 1) + 2^{24k'} \underbrace{c(k' + 1)kn'}_{\text{flow algorithm}}, \quad (2)$$

for some $c > 0$. It is difficult to obtain a good solution, but by induction on k' we get the following loose upper bound.

$$T(n', k') \leq 3c2^{24k'} k'! kn. \quad (3)$$

For $k' = 0$, this bound is valid. For $k' \geq 1$, we have:

$$\begin{aligned} T(n', k') &\leq 24k'T(n' - 1, k' - 1) + c2^{24k'} (k' + 1)kn' \\ &\leq 3c24k'2^{24(k'-1)}(k' - 1)!kn + c2^{24k'} (k' + 1)kn' \quad \text{by induction hypothesis} \\ &= c2^{24k'} kn \left(\frac{3 \cdot 24k'}{2^{24}} + k' + 1 \right) \leq 3c2^{24k'} knk'! \left(\frac{24}{2^{24}} + \frac{k' + 1}{3k'!} \right) \\ &\leq 3c2^{24k'} k'! kn. \end{aligned}$$

Even though, this is not a tight bound, we have $T(n, k) \geq c'24^k k!(n - k)$, which is $2^{\Omega(k \log k)} n$.

$$\begin{aligned} T(n, 0) &\geq c'n \\ T(n, k) &\geq 24kT(n - 1, k - 1) \\ &\geq c'24k24^{k-1}(k - 1)!(n - k) \text{ by inductive hypothesis} \\ &\geq c'24^k k!(n - k) \end{aligned}$$

4 Our Improved Algorithm

In this section, we discuss how to improve Reed's algorithm. The dependence on k in the running time of Reed's algorithm is huge. We decrease this factor significantly to make the algorithm more applicable. We introduce two main modifications. First, we work with a larger cut-off threshold than Reed's $\frac{|V|}{24k}$. Such an improvement can be achieved by replacing the arbitrary $3/4$ bound by $1 - \epsilon$. But even more is possible by arguing about the weights of connected components instead of the weights of the parts of a bipartition.

The second improvement is to avoid branching on whether there is a representative going into the separator or not. Reed branches on these two cases at the beginning, while we branch 3-fold for every representative.

Note that the second improvement is not obvious. Reed had a good reason to avoid 3-fold branching. If a representative of tree is put into the separator, then, we lose control of the unbounded set of additional vertices of the tree.

4.1 Relax the balancing requirement

Reed's argument starts with a weakly balanced separation by volume (i.e. according to V) that is known to exist. The larger side has at most $2/3$ of the volume, but it might have up to $2/3 + 1/24$ of the weight. The algorithmic split by this weight partition might pick a set with another $1/24$ fraction more volume. Thus the worst kind of volume split found is now $3/4$ to $1/4$. Recall that these differences are bounded this way for the following reason. When the weight carrying root

of a tree is on one side, some of its small subtrees rooted at the children can be partially on the other side. But the separator prevents more than k small subtrees to have any part on a different side than the root, and each small subtree contains less than $\frac{|V|}{24k}$ vertices. Instead of $3/4$, one can choose any number strictly between $2/3$ and 1 . If $1 - \epsilon$ is chosen, then the constant 24 is replaced by $\frac{1}{((1-\epsilon)-2/3)/2} = \frac{6}{1-3\epsilon} \leq 6 + 24\epsilon$ for $\epsilon \leq 1/12$.

A better improvement is possible by a modification of the analysis. If G is connected, then we know that a balanced V -separator S' exists. It is a separator by volume. We fix such an S' for the analysis. The algorithm does not have to find it. S' defines a set of connected components G_1, \dots, G_t of $G - S'$ with vertex sets V_1, \dots, V_t . We focus on the set \mathcal{B} of those branches of the algorithm which put any vertex $v \in W'$ into the separator if and only if $v \in S'$. For each branch of \mathcal{B} , the weight of G_i is the weight of $W' \cap V_i$. We choose a larger value of s , namely $s = \lfloor (1/4 - \epsilon/2) \frac{|V|}{k} \rfloor$. Because $|V_i| \leq \frac{|V|}{2}$, the weight of each connected component G_i is at most $(3/4 - \epsilon/2)|V|$. We group V_1, \dots, V_t into L' and R' such that neither L' nor R' have weight more than $(3/4 - \epsilon/2)|V|$.

Now, there exists exactly one branch of \mathcal{B} that places a vertex $v \in W'$ left if and only if $v \in L'$. Otherwise such a vertex $v \in W'$ is placed right on this branch. Again, note that the algorithm does not know L' and R' . However, as it tries all possible placements of W' , one branch is good.

The algorithm finds a separator S that allows the same 2-partition of $W' \setminus S'$ as S' . The separator S found by the algorithm might be different from S' . However, as S' and S agree on W' , the volume ratio is not worse than $(1 - \epsilon)$ to ϵ .

With this improvement, the constant 24 of Reed's algorithm is replaced by

$$\frac{1}{((1 - \epsilon) - 1/2)/2} = \frac{4}{1 - 2\epsilon} \leq 4 + 12\epsilon \text{ for } \epsilon \leq 1/6.$$

4.2 Main Improvement

The other improvement is to allow the representatives (the roots of the subtrees) to go either, left, right, or into the separator. Once a representative v goes into the separator, we change its weight to 0 . We also delete v from G and unmark all of the vertices in its subtree, so that they can be searched again. Then, we continue the DFS from the parent of v . There is one complication here that we have to take care of; what happens if $G - v$ gets disconnected?

Three cases might happen. We cover all the cases and show that they can be handled. Let x, y , and z be three types of children of v with the small subtrees (with size $< \frac{n}{Ck}$) τ_x, τ_y , and τ_z , rooted respectively (Fig. 3(a)). The cases are:

- There is a back edge from a vertex in τ_x to an ancestor of v . If we delete v from the tree, the vertices of τ_x will still be searched because they are connected to an ancestor of v . So, we just need to color them white once more. Hence, we do not need to worry about this case.
- There is a vertex p in τ_y with subtree τ_q attached to it with q as its root such that q is a representative below τ_y which has not gone into the separator (this is a bottom-up approach). Even though this case seems to be troublesome, we can fix it. Let $\tau_{p,y}$ be τ_y rooted at p (dangling from p). Make p a child of q . Now, the problem has been fixed (Fig. 3(b)). The same reasoning applies when v is a root and its subtree is too small.

Notice that in this case, it is possible that there is a back edge to some vertex higher up as well, but it does not change anything here. We handle it by attaching to the subtree below it. Alternatively, one could handle it like the previous case. Both are fine.

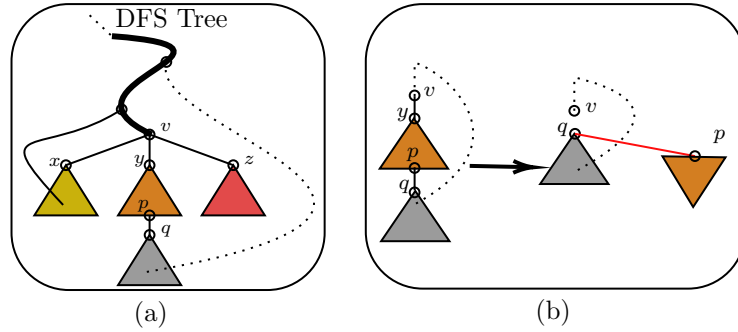


Figure 3: How to fix the cases after a representative goes to the separator. Here, v and q are representatives and v is going to be sent into the separator.

- There is no back edge from τ_z to an ancestor of v , and there is no subtree below. So deleting v makes τ_z disconnected from the entire tree, and this is only for our advantage.

Now, the question is how to distinguish between these three cases?

We just need to distinguish between case 2 and the other two cases (together) because case 1 and 3 are handled the same way.

Cases 1 and 3 are actually the simple ones. We just need to color the vertices in their subtrees white and then we are done. So, we have to find a way to recognize case 2. For that, each vertex u in $\tau(y)$ saves a possible representative q with an edge to some $p \in \tau(y)$ such that p is an ancestor of q (defined as *rep_connected* in the pseudocode). If no such representative exists, simply set its value to nil. Notice that we just need to store one representative connected down below since it is sufficient to handle this case as we just described. Each vertex passes that information to its parent. So, if $rep_connected[y] \neq nil$, we know that the case 2 has happened. Otherwise, case 1 or 3 has happened and we do not need to know which one since they are handled the same way. Now, let's go back the algorithm description.

The main difference here is that Reed branches in the beginning and considers two cases. In the first case, none of the representatives goes to the separator, and in the second case at least one goes to the separator. In the second case, Reed's algorithm branches into at most $24k$ (upper bound for the number of subtrees). We want to avoid these branches and each time only branch into three cases. Assume we want to decide where to put v (a representative with weight $w(v)$). Let L , S , and R be the left, the separator, and the right sets, respectively. If we put v into L (or R), usually most of the vertices in its subtree will be in L (or R) as well. In case v goes to S , we release the other vertices of its tree to be searched again (as we described above and also as it is mentioned in the pseudocode). We have to mention that unlike Reed, we do not decide at the beginning if at least one vertex is going to the separator. Instead, we consider this case for every representative only when we handle that representative (remember that we handle the representative in a serial way).

Theorem 2 *Let $C_0 = 4 + 4 \log 1.25 + \log 5 < 7.60965$. For any $C > C_0$, there exists a 5-approximation algorithm that solves treewidth in time $\mathcal{O}(2^{Ck} n \log n)$, where k is treewidth+1.*

Proof:

Here, we include the pseudocodes of our algorithms (4 and 5), but before that we introduce the global variables.

Global Variables

- G : the input graph
- $root$: a fixed vertex to start all the DFS visits
- W' : the list containing all the representatives (that do not go into the separator)
- \mathcal{W} : the list containing the weight of every representative
- $decision$: the string that dictates where we should place the representatives.
- S : the separator
- k : an upper bound on the capacity of the separator in the beginning.
- k' : an upper bound on the capacity of the separator at any given time.

The following section includes additional explanations on Algorithms 4 and 5 as well the proof of correctness.

4.3 The Correctness of Our Algorithm

We use a subroutine of Robertson-Seymour’s 4-approximation algorithm [18,20] to find a balanced separator of any given set W . We call this splitting “split by W ”. If we continue just by splitting W , we end up having quadratic dependence on n in the running time in a worst case scenario. In order to avoid this, after every $\log_{\frac{3}{2}} k$ such splittings, we split based on V (the reasoning can be found in Section 3.4). Splitting by V is handled by finding a collection of representatives and then splitting the representatives in a balanced way. As we mentioned earlier, at most $\frac{n}{C_\epsilon}$ vertices would be on the opposite side of their representatives. Hence, this way of splitting procedure splits the entire graph in a somewhat balanced way without actually deciding on every vertex (if there exists a balanced separator at all). Below we explain the latter part in more detail.

We do a DFS and the first time that we post-visit a vertex with at least $\frac{n}{C_\epsilon k}$ vertices in its subtree, we call that vertex a representative which represents the vertices of its subtree. Then, we decide what to do with that representative.

Algorithm 4 initiates the process by calling the DFS on G starting from a fixed vertex $root$. The reason that we fix the starting point is that we want to have a deterministic DFS search so that when we change our “ $decision$ ”, the DFS follows the same order.

$Decision$ is a string $\in \{0,1\}^*$ that tells us what to do after we find a representative. If one $decision$ was not good (i.e., it did not give us a balanced separator), we check the next $decision$.

We start with the decision that no representative should go into the separator ($decision = ""$). Notice that $decision$ is a string. The reason is that we want to keep the leading 0’s. Decision either is empty or ends with a 1 (we stop at the right-most 1 and assume the string is padded with an unlimited⁵ number of 0s to the right). If necessary, we keep adding those 0’s to the right.

If $decision[i] = 0$, we send the i -th representative left/right. If it was 1, we send that vertex into the separator and do not consider it a representative anymore. If $|decision|$ is too short, this means it was supposed to be 0 (because the string can be padded with unlimited 0s). Then, we add 0 to the R.H.S. of the $decision$ string and send this representative left/right. Sending left/right

⁵In fact, $C_\epsilon k$ is sufficient.

Algorithm 4: The procedure to initialize the DFS

```

1 Procedure INIT_DFS()
2 Declarations:
3 *  $\mathcal{R}$ .tree_info = (list of vertices  $\mathcal{R}$ .vertices, vertex rep_connected)
4 *  $s = \lfloor (1/4 - \epsilon/2) \frac{|V|}{k} \rfloor$  //  $s$ : the size bound for splitting off an  $\mathcal{R}$ -tree.
5 * decision = "" // No representative should go into the separator
6 * while TRUE do
7   * Initialize  $W'$  and  $\mathcal{W}$  with empty lists //  $W'$  is the (ordered) list of the
     representatives and  $\mathcal{W}$  is the (ordered) list of their corresponding
     weights.  $\mathcal{W} = (w[i]_{1 \leq i \leq |W'|})$ , where  $w[i]$  is the weight of the  $i$ -th
     representative.
8   *  $k' = k$ 
9   * for all  $v \in V$  do
10    | * color[ $v$ ] = WHITE
11  end
12  * vertex_rep = DFS_VISIT(root)
13  *  $x.\mathcal{R}$ .vertices =  $x.\mathcal{R}$ .vertices  $\circ$  vertex_rep. $\mathcal{R}$ .vertices, where  $x$  is the last
     representative added to  $W'$  // If the topmost subtree has smaller size than
      $s$ , we hook it onto the last representative  $x$ .
14  * if  $x.rep\_connected == nil$  then
15    | *  $x.rep\_connected = vertex\_rep.rep\_connected$ 
16  end
17  * if  $k' == 0$  then
18    | * if connected components of  $G - S$  can form a balanced bipartition  $(L, R)$  then
19      | | * return  $(L, R)$ 
20    | end
21  end
22  * else
23    | try all partitions of  $W'$  into  $X$  and  $Y$  to see if there exists a balanced separator.
24    | * if such a partition exists then
25      | | * return  $(L, R)$  //  $L$  is  $X$  along with all the vertices represented
26      | | by vertices of  $X$  (analogously  $R$  corresponds to  $Y$ ).
27    | end
28  end
29  // Now, we update the decision to look for the next possibility
30  * decision = lexicographically next string of at most the same length. // For
     instance, 0011001 follows 011000 and 101 follows 1001111
31  * if no such string exists then
32    | * return "Treewidth is  $> k$ "
33  end
34 End Procedure

```

Algorithm 5: Main recursive procedure of DFS-Trees

Result: \mathcal{R} tree_info about the current partial \mathcal{R} tree

```

1 Procedure DFS-VISIT( $u$ )
2 *  $this = \text{new } \mathcal{R}\text{-tree\_info}(\{u\}, \text{nil})$ 
3 *  $\text{color}[u] = \text{GRAY}$ 
4 * while there exists  $v$  adjacent to  $u$  such that  $\text{color}[v] == \text{WHITE}$  do
5   | *  $\text{child} = \text{DFS\_VISIT}(v)$ 
6   | *  $\text{this.}\mathcal{R}\text{-vertices} = \text{this.}\mathcal{R}\text{-vertices} \circ \text{child.}\mathcal{R}\text{-vertices}$ 
7   | *  $\text{color}[v] = \text{BLACK}$ 
8   | * if  $\text{this.rep\_connected} == \text{nil}$  then
9     |   | *  $\text{this.rep\_connected} = \text{child.rep\_connected}$  // Each vertex remembers one of
9     |   |   the possible representatives to which one of its descendants in the
9     |   |   current  $\mathcal{R}$ tree is connected to and passes it to its parent.
10  | end
11 end
12 *  $\text{length} = |W'| + (k - k')$ 
13 * if  $|\text{this.}\mathcal{R}\text{-vertices}| \geq s$  then
14   | * if  $\text{len}(\text{decision}) < \text{length}$  // 0 means place it left/right. Also, we assume
14   |   that the decision is padded with unlimited zeros to the right of
14   |   decision that we do not show after the last 1, where 1 means place into
14   |   the separator
15   | then
16   |   | *  $\text{decision} = \text{decision} \circ "0"$  // where  $\circ$  is string concatenation
17   | end
18   | * if  $\text{decision}[\text{length}] == 0$  // "0" means place it left/right
19   | then
20   |   | *  $W' = W' \cup \{this\}$ 
21   |   | * return  $(\emptyset, u)$ 
22   | end
23   | * else
23   |   | // i.e.  $\text{decision}[\text{length}] == 1$  which means send  $this$  into the
23   |   | separator
24   |   | *  $k' = k' - 1$ 
25   |   | * if  $\text{this.rep\_connected} \neq \text{nil}$  // Handling the middle case of Figure 3
26   |   | then
27   |   |   | *  $w[\text{this.rep\_connected}] = w[\text{this.rep\_connected}] + |\text{this.}\mathcal{R}\text{-vertices}|$  * return
27   |   |   |  $this$ 
28   |   | end
29   |   | * else
30   |   |   | for all  $z \in \text{this.}\mathcal{R}\text{nodes}$  do
31   |   |   |   | *  $\text{color}[z] = \text{WHITE}$  // Either the left or the right case of
31   |   |   |   |   Figure 3 has happened. In either case we do not need to
31   |   |   |   |   worry.
32   |   |   | end
33   |   |   | * return  $(\emptyset, \text{nil})$ 
34   |   | end
35   | end
36 end
37 End Procedure

```

means just place it in W' , for now. At the end of this iteration, check all the possible ways of splitting W' into X and Y to find a balanced separator. If such a split exist, return (L, R) , where $L = X \cup \{\text{any vertex represented by a representative in } X\}$ (analogously R corresponds to Y). If none exist, check the next *decision* and start over.

Then, the while loop on line 6 checks the entire search space. For any fixed *decision*, we deterministically know what should happen to the representatives that we find. Each time we have a new decision, we set the variables back to their original values and restart searching for a balanced separator.

In line 12, we call the DFS_Visit on *root*. DFS_Visit returns a set of vertices (either the subtree that is too small to form a group represented by a representative or \emptyset if they form a group) and the information whether any vertex in the set returned in the first argument is connected to another representative down below. (We will see later why we need this.) We refer to this pair of information as \mathcal{R} _tree_info (check line 3 of Algorithm 4). There is a chance that the last vertex post-visited (which is *root*) is not a representative, meaning that the size of its subtree is smaller than the threshold s . In this case, we hook the subtree rooted at *root* to the last representative that we have found (similar to part b of Figure 3). This happens in line 13.

The main process happens in DFS_VISIT. Whenever it post-visits a vertex, it checks whether its subtree is big enough. If the size of the current subtree hits the threshold, it forms a group and decides whether the current representative should go left/right or into the separator based on *decision* (0 means send it into the separator and 1 means send it left/right). Otherwise, it passes the set of nodes to the parent along with a possible representative connected to its subtree down below.

In line 28 of Algorithm 4, we update the decision if the current decision has failed.

The next *decision* is the next string in lexicographical order with at most the same length. For example

$$\left\{ \begin{array}{ll} 1 & \text{follows } 0 \\ 100111 & \text{follows } 100110111 \\ 001011 & \text{follows } 00101011111111 \end{array} \right.$$

The reason that we do not check all the decision space is that if *decision* leads us to a failure, there is no reason to check a string with *decision* as one of its prefixes. Notice that every time we run out of the characters on the R.H.S. of *decision*, we just concatenate a '0' to the R.H.S. of it.

If $tw(G) \leq k$, then there exists a tree decomposition of G namely \mathcal{T} with width (at most) k . Then, based on Lemma 2, there exists a strong centroid x in \mathcal{T} . Both \mathcal{T} and x are unknown, but they do exist (if we knew, we already had a tree decomposition of width k and a balanced separator). The connected components of $G \setminus B_x$ can form a bipartition (L, R) such that no part has more than $\frac{2}{3}n$ vertices. For the sake of argument, fix \mathcal{T} , x , L , and R . Every $W' \subseteq V$ is partitioned into parts in L (name it X), parts in separator $S = B_x$, and parts in R (name it Y). The weight of each part is at most $\left(\frac{2}{3} + \frac{1}{C_\epsilon}\right)n$. One of the branches picks W' . On the branch that picks W' , the separator B_x lies between X and Y . There is no guarantee that we find this separator, but the algorithm finds at least one separator of size $\leq k$ separating X from Y in a balanced way. Once more, up to $\frac{1}{C_\epsilon}n$ vertices are on the opposite side of their representative. Now, the larger side has at most $\frac{2}{3}n + \frac{2}{C_\epsilon}n$ vertices and the smaller side has at least $\frac{1}{3}n - \frac{2}{C_\epsilon}n$ vertices. As we argued, we set $C_\epsilon \leftarrow \frac{4}{1-2\epsilon}$. This means that in a worst case scenario, our algorithm finds an $(\epsilon, 1-\epsilon)$ -separation. Each part has a constant fraction of the vertices. Hence, we have a somewhat balanced separator.

The approximation ratio analysis is similar to what we described in Section 3.4. We start with $|W| \leq 4k$. It can be written as $3k + k$. Let's call the second term, k_{excess} . Initially, $k_{excess} = k$. Each time we split W , the separator separates W into two parts with the largest part having at most $2/3|W|$ vertices of W . We include the separator itself on both sides before recursing. This means

$$|W| \leq \frac{2}{3}(3k + k_{excess}) + \underbrace{k}_{\text{upper bound on the size of the separator}} = 3k + \overbrace{\frac{2}{3}k_{excess}}^{\text{new } k_{excess}}.$$

Each time, the bound on k_{excess} decreases by a factor of $2/3$. After $\log_{3/2} k$ iterations, k_{excess} drops to zero. At this point, we do one split by V . We have to add the separator to both parts and this means $|W|$, which was upper bounded by $3k$, might go up to $4k$ once more. Now, we can do $\log_{3/2} k$ splits by W again and continue as described.

Notice that throughout the entire process $|W| \leq 4k$. Ultimately, in order to merge the two tree decompositions that we found for two subproblems, we add the separator to both sides. This means the largest bag can have up to $5k$ vertices and hence, the approximation ratio is 5.

5 Running Time of Our Algorithm

Let G , k , and t be the initial graph, the bound on the size of the separator, and the number of representatives, respectively. t is at most $C_e k$ since the cut-off threshold for the volume of the subtrees was $\frac{n}{C_e k}$. While proceeding with the algorithm at each step, let G' , k' , and t' be the current graph, the current bound on the size (capacity) of the separator, and a bound on the number of representatives still to find, respectively. Each time we send some vertex to the separator, we decrease the capacity by one. We assume a worst case, where the number of \mathcal{R} -trees is t , and the separator found has size k (other cases only speed up the running time). Then, the recurrence for the running time to find a separator of size at most k in G is:

$$T(t, k) \leq T\left(\left\lceil t \cdot \frac{k-1}{k} \right\rceil, k-1\right) + 2T(t-1, k) + Qkn + \mathcal{O}(1), \text{ for } t, k > 0, \tag{4}$$

where Q is the constant factor of the DFS. On the R.H.S., the first term handles the case where the representative goes to the separator. Therefore, k decreases by 1, and the number of subtrees becomes at most $t \cdot \frac{k-1}{k}$ (delete that vertex and continue the DFS). The second term handles the case that the current representative does not go into the separator but left or right. Here, the capacity of the separator is unchanged but the number of subtrees decreases by 1. The third term is the upper bound of the exact running time of the DFS. And the last term $\mathcal{O}(1)$ is the overhead to make the calls. The base cases of the recurrence:

$$\begin{aligned} T(0, k) &\leq Q(k+1)kn && \text{to do } k+1 \text{ DFSs. We find } k \text{ paths from } X \text{ to } Y, \\ &&& \text{and search for augmenting paths.} \\ T(t, 0) &\leq Qkn && \text{We have to test whether } S \text{ is indeed a separator.} \end{aligned}$$

The recurrence in 4 might seem hopeless, so we simplify it (generous bound).

$$T(t, k) \leq T(t, k-1) + 2T(t-1, k) + Qkn + \mathcal{O}(1), \text{ } t, k > 0 \tag{5}$$

Now, we have to solve this recurrence. Our recursion tree starts from the root $T(t, k)$ and has two children $T(t, k - 1)$ and $T(t - 1, k)$, left and right, respectively. This is an unbalanced binary tree. Each strand terminates when one of the arguments of $T(\cdot, \cdot)$ becomes zero. Each time we choose the left branch (putting one representative into the separator), we decrease k by 1. Otherwise (putting the representative and its subtree to the right or left set of the separator), we decrease t by 1 and multiply the value by 2. Let $\#(t - i, 0)$ be the (worst case) number of leaves with the first argument $t - i$ and the second argument $k = 0$, for $0 \leq i < t$ (analogous notation for $\#(0, k - j)$ for $0 \leq j < k$). Observe that $\#(t - i, 0) = \binom{k+i}{i}$, and $\#(0, k - j) = \binom{t+j}{j}$. The first two terms of Equation 5 can be computed at the leaves and the other two terms are spent in every vertex of the recursion tree. The reason is that for the first two terms, we need the results of the children and recursively everything relies on the results of the leaves. Each time we decide to consider a subproblem (any internal node corresponds to a subproblem), we have to find the actual separator. This step takes Qkn time. Also, in order to make the recursive calls, we spend $\mathcal{O}(1)$ time.

Now, we compute the first part (the first two terms).

$$\begin{aligned}
 & \sum_{i=0}^{t-1} \left(\#(t-i, 0) \underbrace{2^i}_{i \text{ right branches}} T(t-i, 0) \right) + \sum_{j=0}^{k-1} \left(\#(0, k-j) \underbrace{2^t}_{t \text{ right branches}} T(0, k-j) \right) \\
 & \leq \sum_{i=0}^t \left(\binom{k+i}{i} 2^i \underbrace{Qk^2n}_{\text{Eq. 5}} \right) + \sum_{j=0}^{k-1} \left(\binom{t+j}{j} 2^t \underbrace{Qkn}_{\text{Eq. 5}} \right) \\
 & \leq Qk^2n \sum_{i=0}^t \binom{k+i}{k} 2^i + Qkn2^t \sum_{j=0}^{k-1} \binom{t+j}{j} \\
 & < Qk^2n \sum_{i=0}^t \binom{k+t}{k} 2^i + Qkn2^t \binom{k+t}{k-1} \\
 & \leq Qk^2n \binom{k+t}{k} (2^{t+1} - 1) + Qk^2n2^t \binom{k+t}{k}
 \end{aligned}$$

Now, we have to compute the second part of Equation 5 where we should look at every internal vertex of the tree. We have $\binom{t+k}{k} - 1$ internal vertices (in worst-case), and in each vertex with value k' , we spend at most $Qk'kn + \mathcal{O}(1) \leq Qk^2n(1 + \mathcal{O}(\frac{1}{n}))$. Hence,

$$\begin{aligned}
 T(t, k) & \leq (2^{t+1} - 1) Qk^2n \binom{k+t}{k} + Qk^2n(1 + \mathcal{O}(\frac{1}{n})) \binom{t+k}{k} \\
 & = Qk^2n \left(2^{t+1} + \mathcal{O}\left(\frac{1}{n}\right) \right) \binom{k+t}{k}
 \end{aligned} \tag{6}$$

Notice that $T(\cdot, k)$ is monotonic (due to the definition of T) and use the fact that $t \leq C_\epsilon k$. Now, we simplify Equation 6 by bounding $T(t, k)$ with $T(C_\epsilon k, k)$ and using Stirling's approximation.

$$\begin{aligned}
 T(t, k) & \leq T(C_\epsilon k, k) \leq Qk^2n \left(2^{C_\epsilon k+1} + \mathcal{O}\left(\frac{1}{n}\right) \right) \binom{(C_\epsilon + 1)k}{k} \\
 & \leq Qk^{\frac{3}{2}}n \left(2^{C_\epsilon k+1} + \mathcal{O}\left(\frac{1}{n}\right) \right) \sqrt{\frac{C_\epsilon + 1}{2\pi C_\epsilon}} \left(1 + \frac{1}{C_\epsilon} \right)^{C_\epsilon k} (C_\epsilon + 1)^k
 \end{aligned}$$

Now, we compute the running time (T_V), when we split based on V . We assume a worst case split of ϵ to $1 - \epsilon$ and worst case separator of size k .

$$\begin{aligned} T_V(n, k) &= T_V(\epsilon n + k, k) + T_V((1 - \epsilon)n + k, k) + T(t, k) \\ &\leq T_V(\epsilon n + k, k) + T_V((1 - \epsilon)n + k, k) \\ &\quad + Qk^{\frac{3}{2}}n \left(2^{C_\epsilon k + 1} + \mathcal{O}\left(\frac{1}{n}\right) \right) \sqrt{\frac{C_\epsilon + 1}{2\pi C_\epsilon}} \left(1 + \frac{1}{C_\epsilon} \right)^{C_\epsilon k} (C_\epsilon + 1)^k \\ &\leq Qk^{\frac{3}{2}} \left(2^{C_\epsilon k + 1} + \mathcal{O}\left(\frac{1}{n}\right) \right) \sqrt{\frac{C_\epsilon + 1}{2\pi C_\epsilon}} \left(1 + \frac{1}{C_\epsilon} \right)^{C_\epsilon k} (C_\epsilon + 1)^k \frac{1}{\epsilon} n \ln n \end{aligned}$$

In the above equation, we use the corollary to the following lemma (Lemma 5). Note that the reason we add k to both recursive calls is that we add the separator to both subproblems.

Lemma 5 Assume $0 < \epsilon \leq \frac{1}{2}$, $0 < c' \leq c$, $2 \leq k$, and $n_1 + n_2 = n$. Then the recurrence

$$f(n + k) \leq \begin{cases} c'(n + k) & \text{if } n - k \leq 4k \\ f(n_1 + k) + f(n_2 + k) + c(n + k) & \text{otherwise,} \end{cases}$$

where $\frac{1}{2}n \leq n_1 \leq (1 - \epsilon)n$ has a solution with $f(n + k) \leq \frac{c}{\epsilon}n \ln n - ck$, for $n \geq 2k$.

Proof: Case 1: $2k \leq n \leq 4k$. Then $n \geq 4$ implying $\ln n > 1$ and

$$f(n + k) \leq c(n + k) < 2cn - ck < \frac{c}{\epsilon}n \ln n - ck.$$

Case 2: $n \geq 4k$ and $n_2 \geq 2k$.

$$\begin{aligned} f(n + k) &\leq f(n_1 + k) + f(n_2 + k) + c(n + k) \\ &\leq \frac{c}{\epsilon} (n_1 \ln n_1 + (n - n_1) \ln(n - n_1)) - 2ck + c(n + k) \\ &< \frac{c}{\epsilon} \left((1 - \epsilon)n \underbrace{(\ln(1 - \epsilon) + \ln n)}_{< -\epsilon} + \epsilon n (\ln \epsilon + \ln n) \right) + c(n - k) \\ &< \frac{c}{\epsilon} n \ln n - c(1 - \epsilon)n + cn \ln \epsilon + c(n - k) \\ &\leq \frac{c}{\epsilon} n \ln n + \underbrace{(\epsilon + \ln \epsilon)}_{< 0 \text{ for } \epsilon \leq \frac{1}{2}} cn - ck \\ &\leq \frac{c}{\epsilon} n \ln n - ck \end{aligned}$$

Case 3: $n \geq 4k$ and $n_2 = n - n_1 < 2k$.

$$\begin{aligned} f(n + k) &\leq f(n_1 + k) + f(n_2 + k) + c(n + k) \\ &\leq \frac{c}{\epsilon} n_1 \ln n_1 - ck + c'(n_2 + k) + c(n + k) \\ &\leq \frac{c}{\epsilon} (1 - \epsilon)n (\ln(1 - \epsilon) + \ln n) - ck + c'(n_2 + k) + c(n + k) \\ &< \frac{c}{\epsilon} (1 - \epsilon)n (-\epsilon + \ln n) - ck + c'(n_2 + k) + c(n + k) \end{aligned}$$

$$\begin{aligned} &\leq \frac{c}{\epsilon} n \ln n - ck - cn \ln n - c(1 - \epsilon)n + c'(n_2 + k) + c(n + k) \\ &< \frac{c}{\epsilon} n \ln n - ck \end{aligned}$$

The last inequality is true, because $n \geq 4k \geq 8$ implying $\ln n > 2$, and $n + k < 3k \leq \frac{3}{4}n$. □

□

Corollary 1 *Under the conditions of the Lemma 5,*

$$f(n) \leq \frac{c}{\epsilon} n \ln n.$$

Now, the total running time of the algorithm (T_t) is:

$$T_t(n, k) \leq Qk^{\frac{3}{2}} \left(2^{C_\epsilon k+1} + \mathcal{O}\left(\frac{1}{n}\right) \right) \sqrt{\frac{C_\epsilon + 1}{2\pi C_\epsilon}} \left(1 + \frac{1}{C_\epsilon} \right)^{C_\epsilon k} (C_\epsilon + 1)^k \frac{(1 + \log_{\frac{3}{2}} k)}{\epsilon} n \ln n, \quad (7)$$

where it takes $\log_{\frac{3}{2}} k$ steps so that the k_{excess} drops to zero (that is when we need to split by V once more).

As we mentioned in Section 4.1, $C_\epsilon = \frac{4}{1-2\epsilon}$. We plug that into Equation 7.

$$T_t(n, k) \leq Qk^{\frac{3}{2}} \left(2^{\frac{4k}{1-2\epsilon}+1} + \mathcal{O}\left(\frac{1}{n}\right) \right) \sqrt{\frac{5-2\epsilon}{8\pi}} \left(1.25 - \frac{\epsilon}{2} \right)^{\frac{4k}{1-2\epsilon}} \left(\frac{4}{1-2\epsilon} + 1 \right)^k \frac{(1 + \log_{\frac{3}{2}} k)}{\epsilon} n \ln n, \quad (8)$$

In Equation 8, let $\epsilon \rightarrow 0$.

$$\lim_{\epsilon \rightarrow 0} T_t(n, k) \leq \frac{1}{\epsilon} \sqrt{\frac{5}{2\pi}} Qk^{\frac{3}{2}} \left(\log_{\frac{3}{2}} k \right) 2^{(4+4 \log 1.25 + \log 5)k} n \ln n \quad (9)$$

□

□

For instance, by setting $\epsilon = 10^{-2}$, we have

$$T_t(n, k) \leq 90Qk^{1.5} \log_{3/2} k 2^{7.718k} n \ln n,$$

and by setting $\epsilon = 10^{-3}$, we get

$$T_t(n, k) \leq 893Qk^{1.5} \log_{3/2} k 2^{7.621k} n \ln n$$

Corollary 2 *There exists a 5-approximation algorithm for treewidth that runs in time $\mathcal{O}(2^{7.61k} n \log n)$.*

6 Conclusion

In this paper, we have given a detailed analysis of Reed’s treewidth approximation algorithm [17]. We have shown that it runs in time $\mathcal{O}(2^{24k} k! n \log n)$, where k is treewidth+1. Furthermore, we have shown that it is a 5 or 7 approximation, depending on how frequently we split based on V .

Then, we have given our improved algorithm which runs in time $\mathcal{O}(2^{7.61k} n \log n)$, with the same approximation ratio.

Our main goal was to obtain a small coefficient of k in the exponent to make the algorithm more applicable. We think it is still possible to further improve it. Also, trying to come up with a better approximation ratio is worthwhile while maintaining the same dependence on k (a 2-approximation [15] is known but the coefficient of k in the exponent is larger).

Finally, another direction for future work could be giving hardness results on the approximation ratio.

References

- [1] E. Amir. Approximation algorithms for treewidth. *Algorithmica*, 56(4):448–479, 2010. doi:[10.1007/s00453-008-9180-4](https://doi.org/10.1007/s00453-008-9180-4).
- [2] S. Arnborg, D. G. Corneil, and A. Proskurowski. Complexity of finding embeddings in a k -tree. *SIAM Journal on Algebraic Discrete Methods*, 8:277–284, 1987. doi:[10.1137/0608024](https://doi.org/10.1137/0608024).
- [3] S. Arnborg, J. Lagergren, and D. Seese. Problems easy for tree-decomposable graphs extended abstract. In *International Colloquium on Automata, Languages, and Programming*, pages 38–51. Springer, 1988. doi:[10.1007/3-540-19488-6_105](https://doi.org/10.1007/3-540-19488-6_105).
- [4] M. Belbasi and M. Fürer. An improvement of Reed’s treewidth approximation. In *International Workshop on Algorithms and Computation*, pages 166–181. Springer, 2021. doi:[10.1007/978-3-030-68211-8_14](https://doi.org/10.1007/978-3-030-68211-8_14).
- [5] U. Bertele and F. Brioschi. On non-serial dynamic programming. *Journal of Combinatorial Theory, Series A*, 14(2):137–148, 1973. doi:[10.1016/0097-3165\(73\)90016-2](https://doi.org/10.1016/0097-3165(73)90016-2).
- [6] H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on Computing*, 25(6):1305–1317, 1996. doi:[10.1137/S0097539793251219](https://doi.org/10.1137/S0097539793251219).
- [7] H. L. Bodlaender. Discovering treewidth. In *International Conference on Current Trends in Theory and Practice of Computer Science*, pages 1–16. Springer, 2005. doi:[10.1007/978-3-540-30577-4_1](https://doi.org/10.1007/978-3-540-30577-4_1).
- [8] H. L. Bodlaender, P. G. Drange, M. S. Dregi, F. V. Fomin, D. Lokshantov, and M. Pilipczuk. A $\mathcal{O}(c^k n)$ 5-approximation algorithm for treewidth. *SIAM Journal on Computing*, 45(2):317–378, 2016. doi:[10.1137/130947374](https://doi.org/10.1137/130947374).
- [9] B. Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Information and Computation*, 85(1):12–75, 1990. doi:[10.1016/0890-5401\(90\)90043-H](https://doi.org/10.1016/0890-5401(90)90043-H).
- [10] U. Feige, M. Hajiaghayi, and J. R. Lee. Improved approximation algorithms for minimum weight vertex separators. *SIAM Journal on Computing*, 38(2):629–657, 2008. doi:[10.1137/05064299X](https://doi.org/10.1137/05064299X).

- [11] J. Flum and M. Grohe. *Parameterized Complexity Theory (Texts in Theoretical Computer Science. An EATCS Series)*. Springer-Verlag, Berlin, Heidelberg, 2006.
- [12] F. V. Fomin, I. Todinca, and Y. Villanger. Large induced subgraphs via triangulations and CMSO. *SIAM Journal on Computing*, 44(1):54–87, 2015. doi:10.1137/140964801.
- [13] R. Halin. S-functions for graphs. *Journal of Geometry*, 8(1-2):171–186, 1976. doi:10.1007/BF01917434.
- [14] T. Kloks. *Treewidth: Computations and Approximations*. Springer, 1994. doi:10.1007/BFb0045375.
- [15] T. Korhonen. A single-exponential time 2-approximation algorithm for treewidth. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 184–192, 2022. doi:10.1109/FOCS52979.2021.00026.
- [16] J. Lagergren. Efficient parallel algorithms for graphs of bounded tree-width. *Journal of Algorithms*, 20(1):20–44, 1996. doi:10.1006/jagm.1996.0002.
- [17] B. A. Reed. Finding approximate separators and computing tree width quickly. In *Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing*, pages 221–228. ACM, 1992. doi:10.1145/129712.129734.
- [18] B. A. Reed. *Tree Width and Tangles: A New Connectivity Measure and Some Applications*, page 87–162. London Mathematical Society Lecture Note Series. Cambridge University Press, 1997. doi:10.1017/CB09780511662119.006.
- [19] N. Robertson and P. D. Seymour. Graph minors. III. Planar tree-width. *Journal of Combinatorial Theory, Series B*, 36(1):49–64, 1984. doi:10.1016/0095-8956(84)90013-3.
- [20] N. Robertson and P. D. Seymour. Graph minors. XIII. The disjoint paths problem. *Journal of Combinatorial Theory, Series B*, 63(1):65–110, 1995. doi:10.1006/jctb.1995.1006.