



# Space-Efficient Vertex Separators for Treewidth

Frank Kammer<sup>1</sup> · Johannes Meintrup<sup>1</sup> · Andrej Sajenko<sup>1</sup> 

Received: 1 October 2020 / Accepted: 3 April 2022 / Published online: 29 April 2022  
© The Author(s) 2022

## Abstract

For  $n$ -vertex graphs with treewidth  $k = O(n^{1/2-\epsilon})$  and an arbitrary  $\epsilon > 0$ , we present a word-RAM algorithm to compute vertex separators using only  $O(n)$  bits of working memory. As an application of our algorithm, we give an  $O(1)$ -approximation algorithm for tree decomposition. Our algorithm computes a tree decomposition in  $c^k n (\log \log n) \log^* n$  time using  $O(n)$  bits for some constant  $c > 0$ . Together with the result of Banerjee et al. (Proceedings of 21st international conference on computing and combinatorics (COCOON 2015). LNCS, vol 9198, Springer, pp 349–360, 2015. [https://doi.org/10.1007/978-3-319-21398-9\\_28](https://doi.org/10.1007/978-3-319-21398-9_28)) we are able to compute a solution for all monadic-second-order problems (MSO) with  $O(n + \tau(k) \cdot p(\log_p n) \log n)$  bits in  $O(\tau(k) \cdot n^{2+(2/\log p)})$  time where  $k$  is the treewidth of the given graph,  $p$  is some arbitrary parameter with  $2 \leq p \leq n$  and  $\tau$  is some function depending on the MSO formula. We finally use the tree decomposition obtained by our algorithm to solve VERTEX COVER, INDEPENDENT SET, DOMINATING SET, MAXCUT and  $q$ -COLORING by using polynomial time and  $O(n)$  bits as long as the treewidth of the graph is smaller than  $c' \log n$  for some problem dependent constant  $0 < c' < 1$ .

**Keywords** FPT · Tree decomposition · Network flow · Subgraph stack

---

Johannes Meintrup and Andrej Sajenko: Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 379157101.

---

✉ Andrej Sajenko  
Andrej.Sajenko@mni.thm.de

Frank Kammer  
Frank.Kammer@mni.thm.de

Johannes Meintrup  
Johannes.Meintrup@mni.thm.de

<sup>1</sup> THM, University of Applied Sciences Mittelhessen, Giessen, Germany

## 1 Introduction

For solving problems in the context of the ever-growing field of big data we require algorithms and data structures that do not only focus on runtime efficiency, but consider space as an expensive and valuable resource. Some reasons for saving memory are that slower memory in the memory hierarchy has to be used, less cache faults arise and the available memory allows us to run more parallel tasks on a given problem.

As a solution researchers began to provide *space-efficient* algorithms and data structures to solve basic problems like graph-connectivity problems [3, 5, 14, 20, 24], other graph problems [27, 32], memory initialization [26, 33], dictionaries with constant-time operations [12, 15, 23] or graph interfaces [6, 27] space efficiently, i.e., they designed practical algorithms and data-structures that run (almost) as fast as standard solutions for the problem under consideration while using asymptotically less space. Several space-efficient algorithms and data structures mentioned above are implemented in an open source GitHub project [31].

Our model of computation is the word RAM where we assume to have the standard operations to read, write as well as arithmetic operations (addition, subtraction, multiplication and bit shift) take constant time on a word of size  $\Theta(\log n)$  bits where  $n \in \mathbb{N}$  is the size of the input. To measure the total amount of memory that an algorithm requires we distinguish between the *input memory*, i.e., the read-only memory that stores the input, and the *working memory*, i.e., the read-write memory an algorithm additionally occupies during the computation. In contrast, the *streaming model* [35] allows us to access the input only in a purely sequential fashion, and the main goal is to minimize the number of passes over the input. There are several algorithms for NP-hard problems on a streaming model, e.g., a kernelization algorithm due to Fafianie and Kratsch [21]. Elberfeld et al. [19] proved that the problem of obtaining a so-called tree decomposition is contained in LSPACE.

We continue the previous research of Banerjee et al. [4] on space-efficient algorithms on the word RAM for NP-hard problems. However, they assumed that a tree decomposition is already given with the input. With our space efficient construction of said tree decomposition we improve upon their result, as it is no longer required to have a tree decomposition as the input. More on this in Sect. 8.

One approach to solve NP-hard graph problems is to decompose the given graph into a tree decomposition  $(T, B)$  consisting of a tree  $T$  where each node  $w$  of the tree has a bag  $B(w)$  containing a small fraction of the original vertices of the graph. The quality of a tree decomposition is measured by its *width*, i.e., the number of vertices of the largest bag minus 1. The *treewidth* of a graph is the smallest width over all tree decompositions for the graph. Having a tree decomposition  $(T, B)$  of a graph, the problem is solved by first determining the solution size of the problem in a bottom-up traversal on  $T$  and second in a top-down process computing a solution for the whole graph. Treewidth is also a topic in current interdisciplinary research, such as smart contracts using cryptocurrency [13] or computational quantum physics [18], which are fields that often work with big data sets. So it is important to have space efficient algorithms.

Several algorithms are known for computing a tree decomposition. For the following, we assume that the given graphs have  $n$  vertices and treewidth  $k$ . Based on ideas

in [39], Reed [37] showed an algorithm for computing a tree decomposition of width  $O(k)$  in time  $O(c^k n \log n)$  for some constant  $c$ . The idea is to find a separator  $S$  of the given input graph  $G = (V, E)$  such that each of the connected components of  $G[V \setminus S]$  contains a constant fraction of the vertices of  $V$ , roughly speaking. This is repeated recursively for each of the connected components of  $G[V \setminus S]$ . Algorithm 1 gives a pseudo-code description of general separator based tree decomposition algorithms. A detailed description of Reed's algorithm can be found in Sect. 4.

---

**Algorithm 1:** Separator based tree decomposition algorithm

---

**Input:**  $G = (V, E)$ ,  $X = \emptyset$   
**Result:** Outputs a tree decomposition  $(T, B)$  of  $G$

- 1  $S \leftarrow \text{findASeparator}(G)$ ;
- 2  $B \leftarrow S \cup (X \cap V)$ ;
- 3 Initialize a tree  $T$  with a single root vertex  $r$ ;
- 4 Initialize a bag  $B(r) = B$ ;
- 5 **foreach** *Connected component*  $C$  of  $G[V \setminus S]$  **do**
- 6      $G' \leftarrow G[C \cup S]$ ;
- 7      $X' \leftarrow (X \cap C) \cup S$ ;
- 8     Recurse for  $G'$  and  $X'$  to obtain a tree decomposition  $(T', B')$  for  $G'$ ;
- 9     Connect the root  $r'$  of  $T'$  to the root  $r$  of  $T$ ;
- 10    Set  $B(w) = B'(w)$  for all nodes  $w$  of  $T'$ ;
- 11 **end**
- 12 Output  $(T, B)$  as the tree decomposition of  $G$ ;

---

A similar result to ours is due to Izumi and Otachi ([28], Theorem 2). They showed that for a given graph  $G$  with treewidth  $k \leq \sqrt{n}$  there exists an algorithm that obtains a tree decomposition of width  $O(k\sqrt{n} \log n)$  in polynomial time and  $O(k\sqrt{n} \log^2 n)$  bits. A refined analysis shows a runtime of  $\Omega(n^2)$ . In comparison, our algorithm uses  $O(n)$  bits for graphs of treewidth  $O(n^{1/2-\epsilon})$  for an arbitrary  $\epsilon > 0$  and obtains a tree decomposition of width  $O(k)$  with a runtime that is quasi-linear in  $n$  and exponential in  $k$ . Depending on the use-case our result can be viewed as an improvement, in particular in cases where  $O(n)$  bits are available and a constant approximation is needed. In the case of low treewidth and an availability of only  $o(n)$  bits our algorithm is not an option.

Further tree-decomposition algorithms can be found in [2, 8, 9, 22, 34]. The basic strategy of repeated separator searches is the foundation of all of these treewidth approximation algorithms, as mentioned by Bodlaender et al. [11]. Using the same strategy, Bodlaender et al. also presented an algorithm that runs in  $2^{O(k)}n$  time and finds a tree decomposition having width  $5k + 4$ .

To obtain a space-efficient approximation algorithm for treewidth we modify Reed's algorithm. We finally use a hybrid approach, which combines our new algorithm and Bodlaender et al.'s algorithm [11] to find a tree decomposition in  $b^k n (\log^* n) \log \log n$  time for some constant  $b$ . The general idea for the hybrid approach is to use our space-efficient algorithm for treewidth only for constructing the nodes of height at most  $b^k \log \log n$ . For the subgraph induced by the bags of the vertices below a node of

height  $b^k \log \log n$ , we use Bodlaender et al.'s algorithm. The most computationally difficult task of this paper is the computation of the separators with  $O(n)$  bits.

Finding separators requires finding vertex-disjoint paths for which a DFS as a subroutine is usually used. The idea for finding a separator for a graph  $G = (V, E)$  is to use standard network flow techniques to construct all vertex-disjoint paths starting from some source vertex  $s$  and ending at some sink vertex  $t$ . We then get a separator by taking a single vertex from each of these vertex-disjoint paths such that  $s$  and  $t$  are in different connected components of  $G[V \setminus S]$ . This description omits some details specifically for constructing balanced separators as needed by Reed's algorithm, but the general idea is the same. More details can be found in Sects. 3 and 4.

All recent space-efficient DFS require  $\Omega(n)$  bits [3, 14, 20, 24] or  $\Omega(n^2)$  time due to  $\Omega(n)$  separator searches [28]. Thus, our challenge is to compute a separator and subsequently a tree decomposition with  $O(n)$  bits with a running time almost linear in  $n$ .

To compute a separator of size at most  $k$  with  $O(n \log k)$  bits, the idea is to store up to  $k$  vertex disjoint paths by assigning a color  $c \in \{0, \dots, k\}$  to each vertex  $v$  such that we know to which path  $v$  belongs. We also number the vertices along a path  $P$  with 1, 2, 3, 1, 2, 3, etc. as a way to encode the direction of the path. To give an intuition, if we output a path  $P$  stored by our scheme that starts at a vertex  $v$  (numbered by 1) we next output the vertex  $u$  that is part of  $P$  and adjacent to  $v$  and numbered 2. From there we continue the same way with the vertex  $w$  adjacent to  $u$  numbered 3 and from there to the vertex  $x$  numbered 1 and so on. By using the numbering, we can reconstruct a subpath  $P'$  of  $P$  and its direction starting from an arbitrary vertex  $v$  of  $P$ : the subpath continues exactly at the neighbor  $u$  of  $v$  numbered higher than  $v$ , or numbered 1 if  $v$  is numbered 3. Since we want to find separators with only  $O(n)$  bits, we further show that it suffices to store the color information only at every  $\Theta(\log k)$ th vertex. We so manage to find separators of size at most  $k$  with  $O(n + k^2(\log k) \log n)$  bits. If  $k = O(n^{1/2-\epsilon})$  for an arbitrary  $\epsilon > 0$ , we thus use  $O(n)$  bits.

Our solution to find a separator is in particular interesting because previous space-efficient graph-traversal algorithms either reduce the space from  $O(n \log n)$  bits to  $O(n)$ , e.g., depth first search (DFS) or breath first search (BFS) [14, 20, 24], or reduce the space from  $O(m \log n)$  to  $O(m)$ , e.g., Euler partition [27] and cut-vertices [29]. In contrast, we reduce the space for the separator search from  $O((n + m) \log n)$  bits to  $O(n)$  bits for small treewidth  $k$ .

Besides the separator search, many algorithms for treewidth store large subgraphs of the  $n$ -vertex,  $m$ -edge input graph during recursive calls, i.e., they require  $\Omega((n + m) \log n)$  bits. We modify and extend the algorithm presented by Reed with space-efficiency techniques (e.g., store recursive graph instances with the so-called subgraph stack [27]) to present an iterator that allows us to output the bags of a tree decomposition of width  $O(k)$  in an Euler-traversal order using  $O(kn)$  bits in  $c^k n \log^* n$  time for some constant  $c$ . To lower the space bound further, we use the subgraph stack only to store the vertices of the recursive graph instances. For the edges we present a new problem specific solution.

In Sect. 2, we summarize known data structures and algorithms that we use afterwards. Our main result, the computation of  $k$ -vertex disjoint paths is shown in Sect. 3. We sketch Reed's algorithm in Sect. 4, where we also show a space-efficient computa-

**Table 1** The space requirements in bits to run the different parts of the algorithm to compute a tree decomposition (t.d.) with applications for an  $n$ -vertex graph with sufficiently small treewidth  $k$

	Standard	Intermediate goal	Final goal
$k$ vertex-disjoint paths	$\Omega(kn \log n)$	$\Theta(n)$ (Sect. 3)	
Balanced vertex separator	$\Omega(kn \log n)$	$\Theta(n)$ (Sect. 4)	
Subgraph stack	$\Omega(kn \log n)$	$\Theta(kn)$ (Sect. 5)	$\Theta(n)$ (Sect. 6)
Iterator for a t.d.	$\Omega(kn \log n)$	$\Theta(kn)$ (Sect. 5)	$\Theta(n)$ (Sect. 7)
NP-hard problems above	$\Omega(kn \log n)$	$\Theta(n)$ (Sect. 8)	

tion of a balanced vertex separator using  $O(n)$  bits. In Sect. 5 we present an iterator that outputs the bags of a tree decomposition using  $O(kn)$  bits. Following that, in Sect. 6, we show intermediate results necessary to further lower the space bound. In Sect. 7 we combine the results of all previous sections with the approximation algorithm of Bodlaender et al. [11] to obtain an iterator using  $O(n)$  bits for graphs with treewidth  $O(n^{1/2-\epsilon})$ . We conclude the paper by showing in Sect. 8 that our tree decomposition iterator can be used together with the framework of Banerjee et al. [5] to solve arbitrary MSO problems. For NP-hard problems like VERTEX COVER, INDEPENDENT SET, DOMINATING SET, MAXCUT and  $q$ -COLORING we provide an alternative faster framework that uses  $O(n)$  bits on graphs with small treewidth. The Table 1 below summarizes the space bound of the algorithms described in this paper.

## 2 Preliminaries

Let  $G = (V, E)$  be an undirected  $n$ -vertex  $m$ -edge graph. If it is helpful, we consider an edge  $\{u, v\}$  as two directed edges (called *arcs*)  $(u, v)$  and  $(v, u)$ . As usual for graph algorithms we define  $V = \{1, \dots, n\}$ . The degree of a vertex  $v \in V$  is accessed via the  $\text{deg}_G$  function. Sometimes we operate in directed graphs, in which case we access the in- and out-degree of  $v \in V$  through functions  $\text{deg}_G^{\text{in}}, \text{deg}_G^{\text{out}} : V \rightarrow \mathbb{N}$  that returns the number of incoming and outgoing edges of a vertex  $v$ , respectively. Moreover, in this case  $\text{deg}_G = \text{deg}_G^{\text{in}} + \text{deg}_G^{\text{out}}$ . We denote by  $N(v) \subset V$  the neighborhood of  $v$ . For  $A = \{(v, k) \in V \times \mathbb{N} \mid 1 \leq k \leq \text{deg}_G(v)\}$ , let  $\text{head}_G : A \rightarrow V$  be a function such that  $\text{head}_G(v, k)$  returns the  $k$ th neighbor of  $v$ . Intuitively speaking, each vertex has an adjacency array. In the following let  $V' \subset V$  be a set of vertices of a graph  $G = (V, E)$ . A *path*  $P$  of a graph  $G = (V, E)$  is an ordering of  $V'$  such that subsequent vertices in the ordering are connected by an edge. An  $s$ - $t$  path  $P$  is a path where the first vertex in the ordering is  $s$  and the last vertex is  $t$  for some  $s, t \in V'$ . We say the  $s$ - $t$  path  $P$  connects  $s$  and  $t$  and we call  $s$  the *start vertex of  $P$*  and  $t$  the *end vertex of  $P$* . We call the vertices of  $P \setminus \{s, t\}$  the *internal vertices of  $P$* . Given two vertices  $u, v$  part of some path  $P$  we call  $v$  the *successor of  $u$*  exactly if  $v$  is ordered directly after  $u$  in  $P$ . In this case  $u$  is called the *predecessor of  $v$* . In the case that  $v$  is ordered somewhere after  $u$  in  $P$  we say  $v$  is *behind  $u$  on  $P$*  and  $u$  is *before  $v$  on  $P$* . Given a path  $P$  in a graph  $G = (V, E)$ , we call an edge  $\{u, v\}$  a *chord* exactly if  $u, v \in P, u \neq v$  and  $v$  is not a successor or predecessor of  $u$ . We call  $P$  *chordless* if

no such edge exists. Given a path  $P = (u_1, u_2, \dots, u_\ell)$  and some integers  $i, j$  such that  $1 \leq i \leq j \leq \ell$  we call  $P' = (u_i, \dots, u_j)$  a *subpath* of  $P$ .

If space is not a concern, it is customary to store each graph that results from a transformation separately. To save space, we always use the given graph  $G$  and store only some auxiliary information that helps us to implement the following graph interface for a transformed graph.

**Definition 1** (*Graph interface*) A data structure for a graph  $G = (V, E)$  implements the *graph interface (with adjacency arrays)* exactly if it provides the functions  $\text{deg}_G : V \rightarrow \mathbb{N}$ ,  $\text{head}_G : A \rightarrow V$ , where  $A = \{(v, k) \in (V \times \mathbb{N}) \mid 1 \leq k \leq \text{deg}_G(v)\}$ , and gives access to the number  $n$  of vertices and the number  $m$  of edges. If  $G$  is directed,  $\text{deg}_G^{\text{in}}$  and  $\text{deg}_G^{\text{out}}$  are also supported.

Let  $G = (V, E)$  be a graph and let  $V' \subseteq V$ . During our computation, some of our graph interfaces can support  $\text{head}_G(v, k)$  and  $\text{deg}_G(v)$  only for vertices  $v \notin V'$ . For vertices in  $V'$ , we can access their neighbors via adjacency lists, i.e., we can use the functions  $\text{adjfirst} : V' \rightarrow P \cup \{\text{null}\}$ ,  $\text{adjhead} : P \rightarrow V$  and  $\text{adjnext} : P \rightarrow P \cup \{\text{null}\}$  for a set of pointers  $P$  to output the neighbors of a vertex  $v$  as follows:  $p := \text{adjfirst}(v)$ ; while  $(p \neq \text{null})$  { print  $\text{adjhead}(p)$ ;  $p := \text{adjnext}(p)$ ; }. We then say that we have a *graph interface with  $|V'|$  access-restricted vertices*.

In an undirected graph it is common to store an edge at both endpoints, hence, every undirected edge  $\{u, v\}$  is stored as an *arc* (directed edge)  $(u, v)$  at the endpoint  $u$  and as an arc  $(v, u)$  at the endpoint  $v$ . We also use the two space-efficient data structures below.

**Definition 2** (*Subgraph stack* [27, Theorem 4.7]) The *subgraph stack* is a data structure initialized for an  $n$ -vertex  $m$ -edge undirected graph  $G_0 = (V, E)$  that manages a finite list  $(G_0, \dots, G_\ell)$ , called the *client list*, of ordered graphs such that  $G_i$  is a proper subgraph of  $G_{i-1}$  for  $0 < i \leq \ell$ . Each graph in the client list implements the graph interface supporting the access operations ( $\text{head}_G$  and  $\text{deg}_G$ ) in  $O(\log^* n)$  time. Additionally, it provides the following operations:

- $\text{push}(B_V, B_A)$ : Appends a new graph  $G_{\ell+1} = (V_{\ell+1}, E_{\ell+1})$  to the client list.  $B_V$  and  $B_A$  are bit vectors that represent the vertices and arcs contained in the subgraph  $G_{\ell+1}$  respectively, with  $|B_V| = |V_\ell|$  and  $|B_A| = 2|E_\ell|$  (each edge is represented by two arcs). In each of the bit vectors a bit at index  $i$  is set to 1 exactly if the respective vertex or arc is contained in  $G_{\ell+1}$ . The push-operation runs in  $O((|B_V| + |B_A|) \log^* n)$  time and uses  $O(|B_V| + |B_A|)$  bits.
- $\text{pop}$ : Removes  $G_\ell$  from the client list in constant time.
- $\text{toptune}$ : Subsequent accesses to the interface of  $G_\ell$  can run in  $O(1)$  time. Runs in  $O((|V_\ell| + |E_\ell|) \log^* n)$  time and uses  $O(n + m)$  bits additional space.

For the special case where we have a constant  $\epsilon$  with  $0 < \epsilon < 1$  such that  $|V_{i+1}| < \epsilon|V_i|$  and  $|E_{i+1}| < \epsilon|E_i|$  for all  $0 \leq i < \ell$ , the entire subgraph stack uses  $\sum_{i=0}^\infty ((\epsilon_1)^i n + (\epsilon_2)^i m) = O(n + m)$  bits of space. We will use the subgraph stack only in this special case with  $\epsilon = 2/3$ .

Using rank-select data structures with fast word RAM operations [7] one can easily support the following data structure.

**Definition 3** (*Static space allocation*) For  $n \in \mathbb{N}$ , let  $V = \{1, \dots, n\}$  be the universe and  $K \subset V$  a fixed set of keys with  $|K| = O(n/\log b)$  for some  $b \leq n$ . Assume that the keys are stored within an array of  $n$  bits so that a word RAM can read  $K$  in  $O(n/\log n)$  time. Then there is a data structure using  $O(n)$  bits that allows us to read and write  $O(\log b)$  bits in constant time for each  $k \in K$  and the initialization of the data structure can be done in  $O(n/\log n)$  time.

We now formally define a tree decomposition.

**Definition 4** (*Tree decomposition* [38], *bag*) A *tree decomposition* of a graph  $G = (V, E)$  is a pair  $(T, B)$  where  $T = (W, F)$  is a tree and  $B$  is a mapping  $W \rightarrow \{V' \mid V' \subseteq V\}$  such that the following properties hold: (TD1)  $\bigcup_{w \in W} G[B(w)] = G$ , and (TD2) for each vertex  $v \in V$ , the nodes  $w$  with  $v \in B(w)$  induce a subtree in  $T$ . For each node  $w \in W$ ,  $B(w)$  is called the *bag* of  $w$ .

Recall from the introduction that the *width* of a tree decomposition is defined as the number of vertices in a largest bag minus 1 and the *treewidth* of a graph  $G$  as the minimum width among all possible tree decompositions for  $G$ . We subsequently also use the well-known fact that an  $n$ -vertex graph  $G$  with treewidth  $k$  has  $O(kn)$  edges [38]. Another well known property of treewidth is the following. Given an  $n$ -vertex graph  $G$  with treewidth  $k$  any minor  $H = (V', E')$  of  $G$  has treewidth  $k' \leq k$ . This also implies that  $|E'| = O(k|V'|)$ .

Our algorithms use space-efficient BFS and DFS.

**Theorem 1** (BFS [20]) *Given an  $n$ -vertex  $m$ -edge graph  $G$ , there is a BFS on  $G$  that runs in  $O(n + m)$  time and uses  $O(n)$  bits.*

**Theorem 2** (DFS [24]) *Given an  $n$ -vertex  $m$ -edge graph  $G$ , there is a DFS on  $G$  that runs in  $O(m + n \log^* n)$  time and uses  $O(n)$  bits.*

The DFS of Theorem 2 assumes that we provide a graph interface with adjacency arrays. Later we run a DFS on an  $n$ -vertex graph  $G$  with treewidth  $k$ , but can only provide a graph interface with  $O(k)$  access-restricted vertices, i.e.,  $O(k)$  vertices are provided by a list interface.

**Lemma 1** *Assume that a graph interface for an  $n$ -vertex,  $m$ -edge graph  $G$  with  $O(k)$  access-restricted vertices for some  $k \in \mathbb{N}$  is given. Assume further that we can iterate over the adjacency list of an access-restricted vertex  $v$  in  $O(f(n, m))$  time for some function  $f$ , whereas we can access an entry in the adjacency array of another vertex still in  $O(1)$  time. Then, there is a DFS that runs in  $O(k \cdot f(n, m) + m + n \log^* n)$  time and uses  $O(n + k \log n)$  bits.*

**Proof** To get an idea consider first a standard DFS starting from a vertex  $r$ . Basically it visits each vertex of the connected component of  $G$  that contains  $r$  while iterating over the incident edges of each visited vertex once. If the DFS visits a vertex  $w$ , then the current  $r$ - $w$ -path is managed in a stack that stores the status of the current iteration over the incident edges of each vertex on the path, i.e., for each vertex on the path, it stores a pointer to the edges of the path such that the DFS can return to a vertex and proceed


with the next edge to explore another path. The space-efficient DFS (Theorem 2) stores only a small part of the entire DFS stack and uses further stacks with approximations of the index pointers to reconstruct removed information from the stack. The reason for storing only small parts of the stack as well as approximate indices is to bound the space-usage by  $O(n)$  bits. Let  $V'$  be the set of access-restricted vertices with  $|V'| = k$ . For the access restricted vertices  $V'$  we are not able to use this space-saving strategy since we only have a list interface, i.e. we are unable to use indices. Thus, for the vertices of  $V'$  we simply ignore the space restrictions imposed by the storage scheme of the DFS and instead store a pointer with  $\Theta(\log n)$  bits into the adjacency list for each vertex  $v' \in V'$  additionally to the storage scheme already in place for the DFS, just as any regular DFS working with adjacency lists would need. This of course uses extra space and thus results in a DFS using  $O(n + k \log n)$  bits instead of  $O(n)$  bits. Since there is one iteration over the incident edges of each of the  $O(k)$  access-restricted vertices, the running time of the DFS increases to  $O(k \cdot f(n, m) + m + n \log^* n)$  time.  $\square$

We make use of standard network flow techniques in the following chapter. A *flow network*  $(G = (V, E), c, s, t)$  consists of a directed graph  $G = (V, E)$ , a *capacity function*  $c : E \rightarrow \mathbb{R}$  and two vertices  $s, t \in V$  called the *source* and *sink*. The *capacity* of an edge  $e$  is  $c(e)$ . For easier definitions it is assumed without loss of generality that for a flow network it holds  $(v, u) \in E$  for each  $(u, v) \in E$ . If this is not the case, the edge can simply be added with capacity 0. A *pseudo-flow* is a function  $f : E \rightarrow \mathbb{R}$  constructed for a given flow network that satisfies the following constraints:  $f((u, v)) = -f((v, u))$  and  $f((u, v)) \leq c((u, v))$  for each  $(u, v) \in E$ . A pseudo flow  $f$  is called a *flow* if for all  $v \in V \setminus \{s, t\}$  it holds  $\sum_{e \in E_{\text{in}}} f(e) = \sum_{e \in E_{\text{out}}} f(e)$  where  $E_{\text{in}}$  is the set of all edges into  $v$  and  $E_{\text{out}}$  is the set of all edges out of  $v$ . Intuitively, the flow into a vertex  $v$  must equal the flow out  $v$ . The *residual capacity* of an edge  $e$  with respect to a flow  $f$  is defined as  $c_f = c(e) - f(e)$ . For a given flow network  $(G = (V, E), c, s, t)$  and a flow  $f$  we can construct a *residual network*  $(G_f = (V, E_f), c_f, s, t)$  with  $c_f$  being the residual capacity and  $E_f = \{e \in E | c_f(e) > 0\}$ . An *augmenting path* is a path  $P$  in a residual network starting at  $s$  and ending at  $t$  such that  $c_f(e) > 0$  for all  $e$  that are part of  $P$ . Intuitively, the flow  $f$  can be increased via  $P$ . A flow  $f$  is called a *maximum flow* if no augmenting paths exists. We only make use of a special type of flow network, called a *unit flow network*, where each edge has a capacity of 1. Many network flow problems are often constructed for the case where  $s$  and  $t$  are not single vertices, but sets of vertices  $S$  and  $T$ . In this case a modified flow network can be constructed by introducing a vertex  $s$  and vertex  $t$  such that  $s$  has only outgoing edges to all vertices of  $S$  and  $t$  has only incoming edges from all vertices of  $T$ . We use standard network flow techniques described further in the next section.

### 3 Finding $k$ Vertex-Disjoint Paths Using $O(n + k^2(\log k) \log n)$ Bits

To compute  $k$  vertex-disjoint  $s$ - $t$ -paths we modify a standard network-flow technique where a residual network [1] and a standard DFS is used.





**Fig. 1** Vertex-disjoint  $s$ - $t$ -paths with colored boundaries. Some regions are exemplary sketched by green rectangles (Color figure online)

**Lemma 2** (Network-Flow Technique [1]) *Given an integer  $k$  as well as an  $n$ -vertex  $m$ -edge graph  $G = (V, E)$  with  $k$  unknown pairwise internal vertex-disjoint  $s$ - $t$ -paths for some vertices  $s, t \in V$ , there is an algorithm that can compute  $\ell \leq k$  internal pairwise vertex-disjoint  $s$ - $t$ -paths using  $O((n + m) \log n)$  bits by executing  $\ell$  times a depth-first search, i.e., in  $O(\ell(n + m))$  time.*

The well-known network-flow technique increases the size of an initially empty set of internal vertex-disjoint  $s$ - $t$ -paths one by one in rounds. It makes usage of (1) a residual network for extending a set of edge-disjoint paths by another one and (2) a simple reduction that allows us to find vertex-disjoint paths by constructing edge-disjoint paths in a modified graph.

We can neither afford storing any copy of the original graph, a graph transformed nor the exact routing of the paths because it would take  $\Omega(n \log n)$  bits. Instead, we use graph interfaces that allow us access to a modified graph on the fly and we partition the paths in parts belonging to small so-called regions that allow us to store exact path membership on the boundaries and recompute the paths in between the boundaries of the regions on the fly if required. For a sketch of the idea, see Fig. 1. Storing the exact path membership only for the boundaries has the drawback that we cannot recompute the original paths inside a region, but using a deterministic algorithm we manage to compute always the same one. It is important to remark that the network-flow technique above also works even if we change the exact routing of the  $s$ - $t$ -paths (but not their number) before or after each round.

For the rest of this section, let  $G = (V, E)$  be an  $n$ -vertex  $m$ -edge graph with treewidth  $k$ , let  $\mathcal{P}$  be a set of  $\ell = O(k)$  internal pairwise vertex-disjoint  $s$ - $t$ -paths for some vertices  $s, t \in V$  and assume that  $G$  has more than  $\ell$  such paths, but not more than  $O(k)$ . Let  $V'$  be the set that consists of the union of all vertices over all paths in  $\mathcal{P}$ . Note that the restriction to  $O(k)$  is natural, as due to our specific use-case of searching for separators of size  $O(k)$  in the following sections. Additionally, it makes the runtime and space analysis more pleasant to read. Note that generally everything that follows works even if  $\ell = \Omega(k)$ , but the runtime and space analysis will be different.

Before we can present our algorithm, we present two lemmas that consider properties of internal pairwise vertex-disjoint paths as given in  $\mathcal{P}$ . To describe the properties we first define a special path structure (with respect to  $G$ ). A *deadlock cycle* in  $G$  with respect to  $\mathcal{P}$  consists of a sequence  $P_1, \dots, P_{\ell'}$  of paths in  $\mathcal{P}$  for some  $2 \leq \ell' \leq \ell$  such that, for all  $1 \leq i \leq \ell'$ , there is a subpath  $x_i, \dots, y_i$  on every path  $P_i$  and there is an edge  $\{x_i, y_{(i \bmod \ell') + 1}\} \in E$ . We call a deadlock cycle *simple* if every subpath consists of exactly two vertices, and otherwise *extended* (see Fig. 2).





Fig. 2 A simple and an extended deadlock cycle (left and right, respectively) (Color figure online)

Fig. 3 The extended deadlock cycle constructed in the proof of Lemma 3 (Color figure online)



**Lemma 3** Assume that the paths in  $\mathcal{P}$  are chordless and that  $G$  has no extended deadlock cycle with respect to  $\mathcal{P}$ . Any other set  $\mathcal{Q}$  of  $\ell$  internal pairwise vertex-disjoint  $s$ - $t$ -paths in  $G[V']$  uses only vertices  $V'$  of the paths in  $\mathcal{P}$ .

**Proof** Note that each path in  $\mathcal{P}$  uses another neighbor of  $s$  and since the paths are chordless,  $s$  can not have another neighbor in  $V'$ . In other words,  $s$  has exactly  $\ell$  neighbors in  $V'$  and each path in  $\mathcal{Q}$  must use exactly one of these neighbors. This allows us to name the paths in  $\mathcal{P}$  and in  $\mathcal{Q}$  with  $P_1, \dots, P_\ell$  and  $Q_1, \dots, Q_\ell$ , respectively, such that the second vertex of  $P_i$  and  $Q_i$  are the same for all  $i \in \{1, \dots, \ell\}$ . The following description assumes suitable indices of the paths in  $\mathcal{P}$  and is illustrated in Fig. 3. Assume for a contradiction that the lemma is not true. This means that  $P_1$  uses a vertex  $v_1 \in V'$  whereas  $Q_1$  leaves the vertices of  $P_1$  at some vertex  $u$  to avoid using  $v_1$ . Let  $v_1$  be the first such vertex (while traversing  $P_1$  from  $s$  to  $t$ ).


Since the paths are chordless,  $Q_1$  must continue with a vertex  $v_2$  of a path  $P_2$ . We call such a continuation a *jump*. Since the paths are chordless, this means that  $Q_2$  must leave before reaching  $v_2$  and jump to a path  $P_3$ . At the end, there must be a path  $Q_i$  ( $i \neq 1$ ) that jumps to a vertex  $w$  on  $P_1$ . If  $w$  is before  $v_1$  (while walking on  $P_1$  from  $s$  to  $t$ ), then  $Q_1$  must leave  $P_1$  before  $v_1$ ; a contradiction. If  $w$  is behind  $v_1$ , we have an extended deadlock cycle in  $G$  with respect to  $\mathcal{P}$ ; again a contradiction.  $\square$

The next lemma show that, if there is a set  $\mathcal{Q}$  of  $\ell$  pairwise internal vertex disjoint paths that uses the same vertices  $V'$  as  $\mathcal{P}$ , then  $\mathcal{Q}$  has the same nice properties of  $\mathcal{P}$  we later require.

**Lemma 4** Assume that the paths in  $\mathcal{P}$  are chordless and that  $G$  has no extended deadlock cycle with respect to  $\mathcal{P}$ . Any other set  $\mathcal{Q}$  of  $\ell$  internal pairwise vertex-disjoint  $s$ - $t$ -paths in  $G[V']$  is chordless and  $G$  has no extended deadlock cycle with respect to  $\mathcal{Q}$ .

**Proof** For a contradiction assume that  $\mathcal{Q}$  has a path  $Q$  with a chord  $\{u, w\}$  such that some vertex  $v$  is in between  $u$  and  $w$  on the path  $Q$ . This allows us to construct a path  $Q'$  from  $Q$  that uses the edge  $\{u, w\}$  instead of the subpath of  $Q$  between  $u$  and  $w$ , leaving  $v$  unused. We so get  $\ell$  paths  $\{Q'\} \cup (Q \setminus \{Q\})$  that do not use  $v \in V'$ , which is a contradiction to Lemma 3.

Assume now that there is an extended deadlock cycle  $Z$  in  $G$  with respect to  $\mathcal{Q}$  and some vertex  $v$  on  $Z$  such that  $v$  is an internal vertex of one of the subpaths of  $Z$ . By



**Fig. 4** On the left an extended deadlock cycle  $Z$  is shown in red with respect to  $\ell'$  paths in  $Q$ . On the right a rerouting is sketched such that the new path  $Q'$  do not use vertex  $v$  anymore (Color figure online)

removing the common edges of  $Z$  and the paths in  $Q$  from the paths in  $Q$  and adding the remaining edges of  $Z$  to the paths in  $Q$ , we get a set of pairwise vertex-disjoint  $s$ - $t$ -paths  $Q'$  in  $G[V' \setminus \{v\}]$ . An example is sketched in Fig. 4. The existence of  $Q'$  is again a contraction to Lemma 3.  $\square$

We begin with a high-level description of our approach to compute  $\mathcal{P}$  space efficiently. Afterwards, we present the missing details. A chordless  $s$ - $t$ -path  $P$  can be computed by a slightly extended DFS and to identify the path it suffices to store its vertices in an  $O(n)$ -bit array  $A$  (Lemma 5). However, if we want to store  $\ell > 1$  chordless pairwise internally vertex-disjoint  $s$ - $t$ -paths in  $A$ , we cannot distinguish them without further information. For an easier description assume that each path with its vertices is colored with a different color. Using an array of  $n$  fields of  $O(\log \ell)$  bits each we can easily store the coloring of each vertex and thus  $\ell$  chordless paths with  $O(n \log \ell)$  bits in total. However, our goal is to store  $\ell$  paths with  $O(n + f(\ell) \log n)$  bits for some polynomial function  $f$ . We therefore define a so-called *path data scheme* that stores partial information about the  $\ell$  paths using  $O(n)$  bits (Definition 5).

Recall that our graph under consideration has treewidth  $k$ . The idea of our scheme is to select and color a set  $B \subseteq V'$  of vertices along the paths with the property that  $|B| = O(n/\log k)$  and  $G[V' \setminus B]$  consists only of small connected components of  $O(k \log k)$  vertices. We call these connected components *regions*.

We show such a set  $B$  exists if the paths have so-called *good* properties (Definition 7). Once a path data scheme is created we lose the information about the exact routing of the original paths, but are able to realize operations to construct  $\ell$  possibly different paths. These operations are summarized in a so-called *path data structure* (Definition 6). In particular, it allows us to determine the color of a vertex  $v$  on a path  $P$  or to run along the paths by finding the successor or predecessor of  $v$  in  $O(\deg(v) + f(k))$  time for some polynomial  $f$  (Lemma 7). The idea here is to explore the region containing  $v$  and use Lemma 2 to get a fixed set of  $\ell$  paths connecting the equally colored vertices. By Lemma 4, the new paths still have our good properties. However, assuming that we have a set  $\mathcal{P}$  of good paths and some new chordless path  $P^* \notin \mathcal{P}$ , whose computation is shown in Lemma 10,  $\mathcal{P} \cup \{P^*\}$  is not necessarily good. Our approach to make the paths good is to find a rerouting  $R(\mathcal{P} \cup \{P^*\})$  that outputs a set  $\mathcal{P}'$  of good paths where  $|\mathcal{P}'| = |\mathcal{P}| + 1$  (Lemma 11).

Let  $V^*$  be the vertices of  $\mathcal{P} \cup P^*$ .

We realize  $R$  by a mapping  $\vec{r} : V^* \rightarrow \{0, \dots, k\}$  that defines the successor of a vertex  $v$  as a vertex  $u \in N(v)$  of another path with color  $\vec{r}(v)$ . Intuitively speaking, we cut the paths in  $\mathcal{P} \cup \{P^*\}$  in pieces and reconnect them as defined in  $\vec{r}$ .

The rerouting may consist of  $O(n)$  successor information, each of  $O(\log k)$  bits. Therefore, we compute a rerouting by considering only the  $s$ - $p$ -subpath  $P'$  of  $P^*$  for some vertex  $p$ . We compute the rerouting in  $O(\log k)$  batches by moving  $p$  closer to  $t$  with each batch. Each batch consists of  $\Theta(n/\log k)$  successor information (or less for the last batch). Using the rerouting  $\vec{r}$  and path data structures storing  $\mathcal{P}$  and  $P^*$ , we realize a so-called *weak path data structure* for the paths  $R(P \cup \{P'\})$ . We call it a weak path data structure because it can not be used to determine the color of a vertex. After that we use it to compute a new valid path data scheme for the paths  $R(P \cup \{P'\})$  such that we need neither the rerouting  $\vec{r}$  nor  $V^*$  anymore (Lemma 13). However, since  $P'$  is only a subpath of  $P^*$  we repeat the whole process described in this paragraph with another subpath  $P''$  of  $P^*$ , i.e., the subpath of  $P^*$  whose first vertex is adjacent to the last vertex of  $P'$ . Finally, we so get a valid path data structure storing  $|\mathcal{P}| + 1$  good  $s$ - $t$ -paths (Corollary 1). Intuitively, this can be thought of as a sliding window algorithm, where the window is moved to adjacent regions.

We now describe our approach in detail. To store a single  $s$ - $t$ -path we number the internal vertices along the path with 1, 2, 3, 1, 2, 3, etc. and store the numbers of all vertices inside an array  $A$  of  $2n$  bits. For a vertex  $v$  outside the paths, we set  $A[v] = 0$ . We refer to this technique as *path numbering*. By the numbering, we know the direction from  $s$  to  $t$  even if we hit a vertex in the middle of the path. Note that  $A$  completely defines the vertices of the single path.

To follow a path stored in  $A$ , begin at a vertex  $v$  and look for a neighbor  $w$  of  $v$  with  $A[w] = (A[v] \bmod 3) + 1$ . Note that this approach only works if  $v$  has only one such neighbor. To avoid ambiguities when following a path we require that the path is chordless.

**Lemma 5** *Assume that we are given access to a DFS that uses a stack to store a path from  $s$  to the current vertex and that runs in  $f_t(n, m)$  time using  $f_b(n, m)$  bits. Then, there is an algorithm to compute a path numbering for a chordless  $s$ - $t$ -path  $P$  in time  $O((n + m) + f_t(n, m))$  by using  $O(n + f_b(n, m))$  bits.*

**Proof** We call the internal subpath  $P' \subseteq P$  from a vertex  $v_i$  to a vertex  $v_j$  *skippable* exactly if  $P'$  consists of at least 3 vertices and there exists a *chord*, i.e., an edge from  $v_i$  to  $v_j$ . The idea is to construct first an  $s$ - $t$ -path  $P$  in  $G$  with the DFS and, directly after finding  $t$ , we stop the DFS and then pop the vertices from the DFS stack and remove skippable subpaths until we arrive at  $s$ .

Let  $P = (s = v_1, v_2, \dots, v_x = t)$  be the vertices on the DFS stack after reaching  $t$  and let  $M$  be an  $n$ -bit array marking all vertices with a one that are currently on the stack. (We say that a vertex  $v$  is *marked in  $M$*  exactly if  $M[v] = 1$ .) The algorithm moves backwards from  $t$  to  $s$ , i.e., by first popping  $t$ , setting  $v_{i+1} = t$  and then popping the remaining vertices from the DFS stack as follows.

Pop the next vertex  $v_i$  from the stack, unmark  $v_i$  in  $M$ , and let  $v_{i-1}$  be the vertex that is now on top of the DFS stack. We check  $v_i$ 's neighborhood  $N(v_i)$  in  $G$  for all vertices  $u \notin \{v_{i-1}, v_{i+1}\}$  that are marked in  $M$ , mark them in an  $n$ -bit array  $M'$  as chords and store in a variable  $x$  the number of these chords.

The internal subpaths between  $v_i$  and all vertices marked in  $M'$  are skippable. We remove the skippable paths as follows: As long as  $x \neq 0$ , we pop a vertex  $u$  from the DFS stack. If  $u$  is marked in  $M'$ , we unmark  $u$  in  $M'$  and reduce  $x$  by one. If  $x \geq 1$ ,

we unmark  $u$  in  $M$ . When  $x = 0$ , the edge  $\{u, v_i\}$  is the last chord, we set  $v_i := u$  and  $v_{i+1} := v_i$  and continue our algorithm as described in this paragraph until  $v_i = s$ , i.e., we search for the next skippable subpath.

Once  $P$  is chordless, we can output the path as follows. Allocate an array  $A$  with 2 bits for each vertex and initialize all entries with 0. Let  $i = 1$ . Starting from  $u = s$  go to its only neighbor  $w \neq u$  that is marked in  $M$ . Set  $i := (i \bmod 3) + 1$  and  $A[w] := i$ . Unmark  $u$  in  $M$ , set  $u := w$ , and continue assigning the path numbering until  $w = t$ . Finally return  $A$ .

Observe that the DFS runs once to find a path  $P$  from  $s$  to  $t$  and then we only pop the vertices from the stack. During the execution of the DFS we manage membership of vertices using only arrays of  $O(n)$  bits. The construction of  $A$  uses a single run over the path by exploring the neighborhood of each marked vertex. In total our algorithm runs in  $O(n + m)$  time and uses  $O(n)$  bits in addition to the time and space needed by the given DFS.  $\square$

Recall that the path numbering does not uniquely define  $\ell > 1$  chordless vertex-disjoint paths since there can be edges, called *cross edges*, between vertices of two different paths in  $\mathcal{P}$ . Due to the cross edges, the next vertex of a path can be ambiguous (see again Fig. 1). Coloring these vertices would solve the problem, but we may have to many to store their coloring.

Our idea is to select a set  $B \subseteq V'$  of vertices that we call *boundary vertices* such that  $|B| = O(n/\log k)$  holds and by removing  $B$  from  $G[V']$  we get a graph  $G[V' \setminus B]$  that consists of connected components of size  $O(k \log k)$ . In particular,  $B$  contains all vertices with *large degree*, i.e., vertices with a degree strictly greater than  $k \log k$ . Note that a graph with treewidth  $k$  has at most  $kn$  edges, and so we have only  $O(n/\log k)$  vertices of large degree. We store the color of all boundary vertices  $B$ . This allows us to answer their color directly. Recall that a region is a connected component in  $G[V' \setminus B]$ . Note that a boundary vertex can be adjacent to vertices part of several regions. Let  $v \in B$  be a vertex of some path  $P$ . To avoid exploring several regions when searching the predecessor/successor of  $v$  on  $P$ , we additionally store the color of the predecessor/successor  $w \notin \{s, t\}$  of each  $v \in B$ . According to our described setting above, we formally define our scheme.

**Definition 5** (*Path Data Scheme*) A path data scheme for  $\mathcal{P}$  in  $G$  is a triple  $(A, B, C)$  where

- $A$  is an array storing the path numbering of all paths in  $\mathcal{P}$ ,
- $B$  is a set of all boundary vertices with  $|B| = O(n/\log k)$  and  $B$  defines regions of size  $O(k \log k)$ , and
- $C$  stores the color of every vertex in  $B$  and of each of their predecessors and successors.

We realize  $A$  as well as  $B$  as arrays of  $O(n)$  bits, and  $C$  as a  $O(n)$ -bit data structure using static space allocation. Altogether our path data scheme uses  $O(n)$  bits.

A further crucial part of our approach is to (re-)compute (parts of our) paths fast—in particular, we do not want to use a  $k$ -disjoint-path algorithm—but we also want to guarantee that vertices of the same color get connected with a path constructed by

a deterministic network-flow algorithm. In other words, our path data scheme must store the same color for each pair of colored vertices that are connected by our fixed algorithm of Lemma 2. We call a path data scheme that has this property *valid* (with respect to our fixed algorithm).

To motivate the stored information of our path data scheme, we first show how it can be used to realize a path data structure that allows us to answer queries on all vertices  $V'$  and not only a fraction of  $V'$ . Afterwards, we show the computation of a path data scheme.

**Definition 6** (*Path Data Structure*) A path data structure supports the following operations where  $v \in V'$ .

- $\text{prev/next}(v)$ : Returns the predecessor and the successor, respectively, of  $v$ .
- $\text{color}(v)$ : Returns the color  $i$  of a path  $P_i \in \mathcal{P}$  to which  $v$  belongs to.

To realize the path data structure, our idea is first to explore the region  $R$  in  $G[V' \setminus B]$  containing the given vertex  $v$  by using a BFS and second to construct a graph of the vertices and edges visited by the BFS. We partition the visited colored vertices inside  $R$  into two sets  $S'$  (successors of  $B$ ) and  $T'$  (predecessors of  $B$ ), as well as extend  $R$  by two vertices  $s'$  and  $t'$  that are connected to  $S'$  and  $T'$ , respectively. Then we sort the vertices of  $R$  and the respective adjacency arrays by vertex IDs and run the deterministic network-flow algorithm of Lemma 2 to construct always the same fixed set of paths (but not necessarily the original paths). The construction of the paths in a region consisting of a set  $U$  of vertices is summarized in the next lemma, which we use subsequently to support the operations of our path data structure. We use the next lemma also to make a path data scheme valid. We so guarantee that our network-flow algorithm connects equally colored vertices and no  $k$ -disjoint path algorithm is necessary. This is the reason why the following lemma is stated in a generalized manner and does not simply assume that a path data scheme is given.

**Lemma 6** *Assume that we are given the vertices  $U$  of a region in  $G$  as well as two sets  $S', T' \subseteq U$  of all successors and predecessors vertices of vertices on the boundary, respectively. Take  $n' = |U|$ . Then  $O(n'k^2 \log^2 k)$  time and  $O(k^2(\log k) \log n)$  bits suffice to compute paths connecting each vertex in  $S'$  with another vertex in  $T'$ .*

**Proof** We construct a graph  $G' = G[U]$ , add two new vertices  $s'$  and  $t'$  and connect them with the vertices of  $S'$  and  $T'$ , respectively. To structurally get the same graph independent of the permutation of the vertices in the representation of the set  $U$  we sort the vertices and the adjacency arrays of the graph representation for  $G'$ . The details to do that are described in the three subsequent paragraphs. Finally we run the algorithm of Lemma 2 to compute all  $s'$ - $t'$ -paths in the constructed graph  $G'$ . Since  $S'$  and  $T'$  are the endpoints of disjoint subpaths of paths in  $\mathcal{P}$ , we can indeed connect each vertex in  $S'$  with another vertex in  $T'$ .

We now show the construction of  $G'$ . We choose an arbitrary  $v \in U$  and run a BFS in graph  $G[U]$  starting at  $v$  three times. (We use a standard BFS with the restriction that it ignores vertices of  $G$  that are not in  $U$ .) In the first run we count the number  $n' = O(k \log k)$  of explored vertices. Knowing the exact number of vertices and knowing that all explored vertices can have a degree at most  $k \log k$  in  $G[U]$ ,

we allocate an array  $D$  of  $n' + 2$  fields and, for each  $D[i]$  ( $i = 1, \dots, n' + 2$ ), an array of  $\lceil k \log k \rceil$  fields, each of  $\lceil \log(k \log k) \rceil$  bits. We will use  $D$  to store the adjacency arrays for  $G'$  isomorphic to  $G[U]$ . For reasons of performance, we want to use indirect addressing when operating on  $G'$ . We use a bidirectional mapping from  $U$  to  $\{1, \dots, n'\}$  and use vertex names out of  $\{1, \dots, n'\}$  for  $G'$ . More exactly, to translate the vertex names of  $G'$  to the vertex names of  $G$ , we use a translation table  $M : \{1, \dots, n'\} \rightarrow U$ , and we realize the reverse direction  $M^{-1} : U \rightarrow \{1, \dots, n'\}$  by using binary search in  $M$ , which can be done in  $O(\log n')$  time per access. For the table we allocate an array  $M$  of  $n'$  fields, each of  $\lceil \log n \rceil$  bits.

In a second run of the BFS we fill  $M$  with the vertices explored by the BFS and sort  $M$  by the vertex names. In a third run, for each vertex  $v$  explored by the BFS, we determine the neighbors  $u_1, \dots, u_x$  of  $v$  sorted by their IDs and store  $M^{-1}(u_1), \dots, M^{-1}(u_x)$  in  $D[M^{-1}(v)]$ . Thus,  $D$  allows constant time accesses to a graph  $G'$  isomorphic to  $G[U]$ . During the third BFS run we want to compute and store the two sets  $S'$  and  $T'$  (using a standard balanced tree representation). Afterwards, we are able to compute the paths in  $G[U]$  and using the mapping  $M$  translate the vertex IDs in  $G'$  back to vertex IDs in  $G$ .

**Efficiency** We now analyze the space consumption and the runtime of our algorithm. Since  $G$  has treewidth  $k$  and  $G'$  consists of vertices of a region of  $G$ , we can follow that  $G'$  has  $n' = O(k \log k)$  vertices and  $O(n'k)$  edges.  $D$  uses  $O(n' \log n)$  bits to store  $n'$  pointers, each of  $O(\log n)$  bits, that point at adjacency arrays. The adjacency arrays use  $O(n'k \log k)$  bits to store  $O(n'k)$  vertex IDs out of  $\{1, \dots, n'\}$ . The translation table  $M$  and the sets  $S'$  and  $T'$  use  $O(n' \log n)$  bits. The BFS can use  $O(n'k \log k)$  bits. The space requirement of an in-place sorting algorithm is negligible. Lemma 2 runs with  $O((n' + kn') \log n') = O(n'k \log k)$  bits. In total our algorithm uses  $O(n'k \log k) + O(n' \log n) = O(k^2(\log k) \log n)$  bits.

To compute the graph  $G'$  and construct  $M$  we run a BFS. Running the BFS comes with an extra time factor of  $O(\log n')$  to translate a vertex ID of  $G$  to a vertex ID of  $G'$  (i.e., to access values in  $M^{-1}$ ) and costs us  $O(n'k \log n')$  time in total. Since  $G'$  has  $O(n'k)$  edges  $D$  has  $O(n'k)$  non-zero entries and a sorting of  $D$  can be done in  $O(n'k)$  time. Sorting  $M$  can be done in the same time. Finally, Lemma 2 has to be executed once for up to  $O(k \log k)$  paths (every vertex of the region can be part of  $S' \cup T'$ ), which can be done in  $O((n'k)(k \log k) \log k)$  time since  $O(n'k)$  bounds the edges of the region and since we get an extra factor of  $O(\log k)$  by using the translation table  $M$ . In total our algorithms runs in time  $O(n'k^2 \log^2 k)$ .  $\square$

Since the number of vertices in a region is bound by  $O(k \log k)$ , we can use Lemma 6 and the coloring  $C$  to support the operations of the path data structure in the bounds mentioned below.

**Lemma 7** *Given a valid path data scheme  $(A, B, C)$  we can realize `prev` and `next` in time  $O(\deg(v) + k^3 \log^3 k)$  as well as `color`( $v$ ) in  $O(k^3 \log^3 k)$  time. All operations use  $O(k^2(\log k) \log n)$  bits.*

**Proof** In the case that a vertex  $v$  is in  $B$ , we find its `color` and the color of its predecessor and successor in  $C$ . (Vertices neighbored to  $s$  or  $t$  have  $s$  as the predecessor and  $t$  as

the successor, respectively.) Thus, we can answer `prev` and `next` by iterating over  $v$ 's neighborhood and determining the two neighbors  $u, w \in V'$  that are colored the same as  $v$ . By using the numbering  $A$  we know the incoming and outgoing edge of the path through  $v$  and so know which of the vertices  $u$  and  $w$  is the predecessor or successor of  $v$ .

For a vertex  $v \notin B$ , we explore  $v$ 's region in  $G[V' \setminus B]$  by running a BFS in  $G[V']$  with the restriction that we do not visit vertices in  $B$ . We use a balanced heap to store the set  $U$  of visited vertices. Moreover, we partition all colored vertices of  $U$  into the set  $S'$  (of successors of  $B$ ) and the set  $T'$  (of predecessors of  $B$ ) with respect to the information in  $C$ . In detail, if a colored vertex  $u \in U$  has an equally colored neighbor  $w \in B$ , then  $u$  is the successor of  $w$  if  $A[v] = (A[w] \bmod 3) + 1$  holds and the predecessor of  $v$  if  $A[w] = (A[v] \bmod 3) + 1$  holds. Having  $U, S'$  and  $T'$  we call Lemma 6 to get the paths in the region.

Note that a path within a region can be disconnected by vertices with large degree so that we can have more than two vertices of each color in  $S'$  and  $T'$ . E.g., a path may visit  $s_1, t_1, s_2, t_2, s_3, t_3, \dots \in U$  with  $s_1, s_2, s_3 \in S'$  and  $t_1, t_2, t_3 \in T'$ . Since our path data scheme is valid, we can conclude the following: assume that we connect the computed subpaths using their common equally colored boundary vertices. By Lemma 3 each solution has to use all vertices and the network-flow algorithm indeed connects  $s_1$  with  $t_1, s_2$  with  $t_2$ , etc. By Lemma 4 the paths and thus our computed subpaths are chordless and have no extended deadlock cycles in  $G$  with respect to  $\mathcal{P}'$ . Using the paths we can move along the path of  $v$  to a colored vertex and so answer  $v$ 's color and both of  $v$ 's neighbors on the path.

**Efficiency** We now analyze the runtime of our algorithm and its space consumption. Realizing the operations `prev` and `next` for a vertex of  $B$  can be done in  $O(\deg(v))$  time since both colored neighbors have to be found, and their color can be accessed in constant time using  $C$ . The operation `color` runs in constant time.

For the remaining vertices we have to explore the region in  $G[V' \setminus B]$ . Since every region has  $n' = O(k \log k)$  vertices and also treewidth  $k$ , the exploration can be done in linear time per region, i.e., in  $O(n'k)$  time using a standard BFS. Filling  $U, S'$  and  $T'$  requires us to add  $O(n')$  vertices into balanced heaps, which uses  $O(n' \log n')$  time. The execution of Lemma 6 can be done in  $O(n'k^2 \log^2 k)$  time and running along the paths to find the actual color runs in  $O(n')$  time. Thus, in total `prev` and `next` run in time  $O(\deg(v) + k^3 \log^3 k)$  and `color(v)` in time  $O(k^3 \log^3 k)$ .

A BFS to explore a region in  $G[V' \setminus B]$  uses  $O(n'k \log n) = O(k^2(\log k) \log n)$  bits, while  $U, S'$  and  $T'$  use  $O(n' \log n)$  bits in total. The execution of Lemma 6 uses  $O(k^2(\log k) \log n)$  bits, which is also an upper bound for all remaining operations.  $\square$

To be able to compute a storage scheme for  $\ell > 1$  paths,  $\mathcal{P}$  must satisfy certain so-called good properties that we summarize in Definition 7.

**Definition 7 (Good paths)** We call a set  $\mathcal{P}'$  of  $s$ - $t$ -paths *good* if and only if

1. All paths in  $\mathcal{P}'$  are pairwise internal vertex disjoint,
2. Each path in  $\mathcal{P}'$  is chordless, and
3. There is no extended deadlock cycle in  $G$  with respect to  $\mathcal{P}'$ .



We call a path  $P'$  *dirty* with respect to a set of good  $s$ - $t$  paths  $\mathcal{P}$  if extending  $\mathcal{P}$  by  $P'$  would invalidate any of the good properties of  $\mathcal{P}$ . The next lemma shows the computation of a valid path data scheme for one path  $P$ . The computation is straightforward. Since there is only one path, this path is uniquely defined by the path numbering in  $A$  and we can run along it and select every  $k \lceil \log k \rceil$  vertex as a boundary vertex and compute and store the remaining information required for a valid path data scheme. After the lemma we always assume that a chordless path  $P$  is given together with a valid path data scheme and thus a path data structure so that we can easily access the predecessor and successor of a vertex on the path.

**Lemma 8** *Given a path numbering  $A$  for one chordless  $s$ - $t$ -path  $P$ , we can compute a valid path data scheme  $(A, B, C)$  for  $P$  in  $O(n)$  time using  $O(n)$  bits.*

**Proof** Since  $A$  stores a single chordless path  $P$ , we can run along it from  $s$  to  $t$  while adding every  $\lceil k \log k \rceil$ th vertex and every vertex of large degree into an initially empty set  $B$ . We then can easily determine the predecessor and successor. Thus, we can store  $C$  using static space allocation. A path data scheme storing a single path is valid by default since there is no second path to that an algorithm could switch. We so get a path storage scheme storing  $\ell = 1$  good  $s$ - $t$ -path in  $O(n)$  time using  $O(n)$  bits.  $\square$

Assume that, for some  $\ell \in \mathbb{N}$ , we have already computed  $\ell$  good  $s$ - $t$ -paths, which are stored in a valid path data scheme, and that there are more than  $\ell$  such paths in  $G$ .

Our approach to compute an  $(\ell + 1)$ th  $s$ - $t$ -path is based on the well-known network-flow technique [1] described in the beginning of the section, which we now modify to make it space efficient.

By the standard network-flow technique, we do not search for vertex-disjoint paths in  $G$ . Instead, we modify  $G$  into a directed graph  $G'$  and search for edge-disjoint paths in  $G'$ . To make the approach space efficient, we do not store  $G'$  separately, but we provide a graph interface that realizes  $G'$ .

**Lemma 9** *Given an  $n$ -vertex  $m$ -edge graph  $G = (V, E)$  there is a graph interface representing an directed  $n'$ -vertex  $m'$ -edge graph  $G' = (V', E')$  where  $V' = \{v', v'' \mid v \in V\}$  and  $E' = \{(u'', v'), (v'', u') \mid \{u, v\} \in E\} \cup \{(v', v'') \mid v \in V\}$  and thus  $n' = 2n$  and  $m' = 2m + n$ . The graph interface allows us to access outgoing edges and incoming edges of  $G'$ , respectively, by supporting the operations  $\text{head}_{G'}^{\text{out}}$ ,  $\text{deg}_{G'}^{\text{out}}$  and  $\text{head}_{G'}^{\text{in}}$ ,  $\text{deg}_{G'}^{\text{in}}$ . The graph interface can be initialized in constant time and the operations have an overhead of constant time and  $O(\log n)$  bits.*

**Proof** We now show how to compute the operations for  $G'$  from  $G$  on the fly. For each vertex  $v$  in  $V$ , we define two vertices  $v' = v$  and  $v'' = v + n$  for  $V'$ . By the transformation we can see that every vertex  $v'$  (with  $v' \leq n$ ) has exactly one outgoing edge to  $v'' = v' + n$  and  $\text{deg}_G(v')$  incoming edges from some vertices  $u'' = \text{head}_G(v', j) + n$  ( $j \leq \text{deg}_G(v')$ ). Moreover, every vertex  $v''$  (with  $v'' > n$ ) has  $\text{deg}_G(v'' - n)$  outgoing edges to some vertices  $u' = \text{head}_G(v'' - n, j)$  ( $j \leq \text{deg}_G(v'' - n)$ ) and one incoming edge from  $v' = v'' - n$ . With these information we can provide our stated operations for  $G'$ .

$$\begin{aligned} \text{deg}_{G'}^{\text{out}}(v) &= \begin{cases} 1 & v \leq n \\ \text{deg}_G(v) & v > n \end{cases} & \text{head}_{G'}^{\text{out}}(v, j) &= \begin{cases} v + n & v \leq n \\ \text{head}_G(v - n, j) & v > n \end{cases} \\ \text{deg}_{G'}^{\text{in}}(v) &= \begin{cases} \text{deg}_G(v) & v \leq n \\ 1 & v > n \end{cases} & \text{head}_{G'}^{\text{in}}(v, j) &= \begin{cases} \text{head}_G(v, j) + n & v \leq n \\ v - n & v > n \end{cases} \end{aligned}$$

The operations  $\text{deg}_{G'}^{\text{out}}, \text{deg}_{G'}^{\text{in}}$  and  $\text{head}_{G'}^{\text{out}}, \text{head}_{G'}^{\text{in}}$  have the same asymptotic bounds as  $\text{deg}_G$  and  $\text{head}_G$ , respectively. □

We next show that we can compute an  $(\ell + 1)$ st paths.

**Lemma 10** *Given a valid path data scheme for a set  $\mathcal{P}$  of  $\ell \in \mathbb{N}$  good  $s$ - $t$ -paths in  $G$ ,  $O(n(k + \log^* n)k^3 \log^3 k)$  time and  $O(n + k^2(\log k) \log n)$  bits suffice to compute an array of  $2n$  bits storing a path numbering of an  $(\ell + 1)$ th chordless  $s$ - $t$ -path  $P^*$ , which can be a dirty path with respect to  $\mathcal{P}$ .*


**Proof** To solve the lemma, we adapt the standard network-flow approach as follows. We first use the graph interface of the Lemma 9 to obtain a graph  $G''$  in that we can search for edge-disjoint paths instead of vertex-disjoint paths in  $G$ . We provide the operations `prev`, `next`, `color` and a virtual array  $A$  that gives access to the paths  $\mathcal{P}$  adjusted to  $G'$  by using a simple translation to the corresponding data structure of  $G$ .

According to the construction of  $G'$  a vertex  $v$  in  $G$  is equivalent to two vertices, an in-vertex  $v' = v$  and an out-vertex  $v'' = v + n$  in  $G'$ . Moreover, every in-vertex  $u'$  in  $G$  has a single outgoing edge and this edge points to its out-vertex  $u'' = u' + n$ , which has only outgoing edges that all point to some in-vertices in  $G'$ . Taken our stored  $\ell$  good paths into account a path  $(v_1, v_2, v_3, \dots) \in \mathcal{P}$  in  $G$  translates into a path  $(v'_1, v''_1, v'_2, v''_2, v'_3, v''_3, \dots)$  in  $G'$  and the edges between them must be reversed in  $G'$ .

Next, we want to build the residual network of  $G'$  and  $\mathcal{P}$ . This means that we have to reverse the edges on the paths of  $\mathcal{P}$ . To identify the edges incident to a vertex  $v$ , it seems to be natural to use `prev`( $v$ ) and `next`( $v$ ). Unfortunately, we cannot effort to query these two operations many times to find out which of the incident edges of  $v$  is reversed since  $v$  can be a vertex of large degree and the runtime of both operations depends on  $\text{deg}_{G'}(v)$ . This issue becomes especially important when using a space-efficient DFS that may query `prev`( $v$ ) more than a constant number of times (e.g., by running several restorations of a stack segment including  $v$ ).

To avoid querying `prev`( $v$ ) for any vertex  $v$  of  $G'$ , we present the DFS a graph  $G''$  where  $v$  has as outgoing edges first all outgoing edges of  $v$  in  $G'$  and then all incoming edges of  $v$  in  $G'$ . In detail, we present the DFS a graph  $G''$  with all vertices of  $G'$  and one further vertex  $d$  with no outgoing edges. We also make sure that the DFS has colored  $d$  black (e.g., start the DFS on  $d$  before doing anything else). A sketch of the graph  $G'$  with a blue path, graph  $G''$  and the reversal of the edges in  $G''$  can be seen in Fig. 5. As we see in the next paragraph, we will use  $d$  as a kind of *dead end*.

Every vertex  $v \neq d$  of  $G''$  has  $\text{deg}_{G'}^{\text{out}}(v) + \text{deg}_{G'}^{\text{in}}(v)$  many outgoing edges defined as follows. The heads of the first  $\text{deg}_{G'}^{\text{out}}(v)$  outgoing edges of  $v$  in  $G''$  are the same as



**Fig. 5** From top left to bottom right, a sketch of  $G, G', G''$  and  $G''$  with reversed edges. Each black vertex is the same vertex  $b$ . The blue edges are edges of a path. The dashed lines in the graphs sketches an edge connecting the middle rightmost with the middle leftmost vertex (Color figure online)

in  $G'$  with the exception that we present edge  $(v, d)$  for possibly one outgoing edge  $(v, z)$  that is reversed (i.e.,  $\text{color}(v) = \text{color}(z) \wedge A[z] = (A[v] - 1) \pmod 3$ ). The heads of the next  $\text{deg}_{G'}^{\text{in}}(v)$  incoming edges of  $v$  in  $G''$  are presented as an edge  $(v, d)$  with the exception that we present possibly one incoming edge  $(z, v)$  that is reversed as an outgoing edge  $(v, z)$ .


Intuitively speaking, this ensures that a DFS backtracks immediately after using edges that do not exist in the residual graph. Note that we so can decide the head of an edge in  $G''$  by only calling `color` twice.

The graph  $G''$  has asymptotically the same amount of vertices and edges as  $G$ . To guarantee our space bound we compute a new  $s$ - $t$ -path  $P$  in  $G''$  using a space-efficient DFS on  $G''$  and consecutively make  $P$  chordless by Lemma 5. Finally, we transform  $P$  into a path  $P^*$  with respect to  $\mathcal{P}$  by merging in and out vertices together.

**Efficiency** Concerning the running time we use a space-efficient DFS (Lemma 2) to find a new  $s$ - $t$ -path in  $G''$ . To operate in  $G''$  we need to query the color of vertices. We pay for that with an extra factor of  $O(k^3 \log^3 k)$  (Lemma 7) in the running time. Since our graph has  $O(n)$  vertices and  $O(kn)$  edges we find a new  $s$ - $t$ -path in  $G''$  in time  $O(n(k + \log^* n)k^3 \log^3 k)$ . The time to transform the found path into  $P^*$  and store it can be done in the same bound. The total space used for the algorithm is  $O(n + k^2(\log k) \log n)$  bits— $O(n)$  bits for the space-efficient DFS,  $O(k^2(\log k) \log n)$  bits for Lemma 7, and  $O(n)$  bits for storing  $P^*$  and its transformation.  $\square$

Let us call an  $s$ - $t$ -path  $P^*$  a *clean path* with respect to  $\mathcal{P}$  if  $\mathcal{P} \cup \{P^*\}$  is a set of good paths, and a *dirty path* with respect to  $\mathcal{P}$  if for all  $P \in \mathcal{P}$ , the common vertices and edges of  $P^*$  and  $P$  build subpaths each consisting of at least two vertices and used by the two paths in opposite direction when running from  $s$  to  $t$  over both paths. Intuitively, if we construct an  $(\ell + 1)$ st path  $P^*$  in the residual network of  $G'$  with  $\ell$  paths in  $\mathcal{P}$ , then  $P^*$  can run backwards on paths in  $\mathcal{P}$ , and we call  $P^*$  a dirty path (see Fig. 6). After a rerouting, we obtain  $\ell + 1$  internal vertex disjoint  $s$ - $t$ -paths. To

**Fig. 6** The green subpath of  $P^*$  is clean until it hits a common vertex  $u$  with  $P_3$ , then it is a dirty path (Color figure online)



store the paths  $\mathcal{P} \cup \{P^*\}$  with our path data structure, the paths must be good, which we can guarantee through another rerouting. The details are described below.

Let  $P^*$  be an  $s$ - $t$ -path returned by Lemma 10. We first consider the case where  $P^*$  is a dirty path with respect to  $\mathcal{P}$ . To make the paths good we cut the paths into subpaths and by using the subpaths we construct a set of new  $\ell + 1$  good  $s$ - $t$ -paths. To achieve that we have to get rid of common vertices and extended deadlock cycles. We handle common vertices and extended deadlock cycles in a single process, but we first briefly sketch the standard network-flow technique to remove common vertices of  $P^*$  with a path  $P^c \in \mathcal{P}$ . By the construction of the paths, the common vertices of an induced subpath are ordered on  $P^c$  as  $v_{\sigma_1}, v_{\sigma_2}, \dots, v_{\sigma_x}$  (for some function  $\sigma$  and  $x \geq 2$ ) and on  $P^*$  as  $v_{\sigma_x}, v_{\sigma_{x-1}}, \dots, v_{\sigma_1}$ .  $P^c$  and  $P^*$  can be split into vertex disjoint paths by (1) removing the vertices  $v_{\sigma_2}, \dots, v_{\sigma_{x-1}}$  from  $P^*$  as well as from  $P^c$ , (2) by rerouting path  $P^c$  at  $v_{\sigma_1}$  to follow  $P^*$ , and (3) by rerouting  $P^*$  at  $v_{\sigma_x}$  to follow  $P^c$ .


We denote by  $R(\mathcal{P} \cup \{P^*\})$  a rerouting function that returns a new set of  $|\mathcal{P} \cup \{P^*\}|$  good paths. We so change the successor and predecessor information of some vertices of the paths  $\mathcal{P} \cup \{P^*\}$ . Let  $V^*$  be a set of all vertices in  $R(\mathcal{P} \cup \{P^*\})$ . To define successor and predecessor information for the vertices we use the mappings  $\vec{r}, \overleftarrow{r} : V^* \rightarrow \{0, \dots, k\}$  where for each vertex  $u \in V^*$  with  $\vec{r}(u) \neq 0$ ,  $u$ 's new successor is a vertex  $v \in N(u) \cap V^*$  with  $\vec{r}(u) = \text{color}(v)$ . Similarly we use  $\overleftarrow{r}$  to define a new predecessor of some  $u \in V^*$ . The triple  $(V^*, \vec{r}, \overleftarrow{r})$  realizes our rerouting  $R$ .

To avoid using too much space by storing rerouting information for too many vertices our approach is to compute a rerouting  $R(\mathcal{P} \cup \{P'\})$  only for an  $s$ - $p$ -subpath  $P'$  of  $P^*$  where  $p$  is a vertex of  $P^*$ . This means that  $R(\mathcal{P} \cup \{P'\})$  consists of  $\ell$   $s$ - $t$ -paths and one  $s$ - $p$ -path.


Moreover, we make the rerouting in such a way that, for all paths  $P_i \in R(\mathcal{P} \cup \{P'\})$ , there is a vertex  $v_i$  with the following property: replacing the  $v_i$ - $t$ -subpath of  $P_i$  by a virtual edge  $\{v_i, t\}$  for all paths  $P_i$  we get  $\ell + 1$  good paths. Let  $\mathcal{P}^c$  consist of the  $s$ - $v_i$ -subpaths for each  $P_i$ . We then call the set of vertices part of the paths in  $\mathcal{P}^c$  a clean area  $Q$  for  $R(\mathcal{P} \cup \{P'\})$ .

The idea is to repeatedly compute a rerouting in  $O(\log k)$  batches and thereby extend  $Q$  with each batch by  $\Theta(n/\log k)$  vertices—the last batch may be smaller—such that we store  $O(n/\log k)$  entries in  $\vec{r}$  and  $\overleftarrow{r}$  with each batch. After each batch we free space for the next one by computing a valid path storage scheme storing good  $s$ - $t$ -paths  $R(\mathcal{P} \cup \{P'\}) \setminus P''$ , where  $P''$  is the only path in  $R(\mathcal{P} \cup \{P'\})$  not ending in  $t$ , and a valid path storage scheme for the  $s$ - $t$ -path  $\mathcal{P}^{**}$  obtained from  $P^*$  by replacing  $P'$  via  $P''$ . (For an example consider the  $p$ - $w$ -subpath of  $P^*$  in Fig. 7a, which is replaced by a beginning of  $P_2$  in Fig. 7d.)

We now sketch our ideas of the rerouting that removes extended deadlock cycles in  $\mathcal{P} \cup \{P^*\}$ . Recall that every extended deadlock cycle must contain parts of  $P^*$  since



(a) Look for common vertices or cross edges. (b) Explore subpaths with ext. deadl. cycles.



(c) Run a DFS to find a path from  $u$  to  $w$ . (d) A rerouting over cross edges.

Fig. 7 Steps and data structures of the algorithm to create good paths (Color figure online)

all remaining paths are good. Hence by moving along  $P^*$  we look for an vertex  $u$  of  $P^*$  that is an endpoint of a cross edge  $\{u, v\}$ . Since a deadlock cycle is a cycle, there must be some vertex  $w$  after  $u$  on  $P^*$ , that is connected by subpaths of  $\mathcal{P}$  and cross edges connecting the paths in  $\mathcal{P} \cup \{P^*\}$ .

Since common vertices and deadlock cycles may intersect, vertices  $u$  and  $w$  can have a cross edge or be common vertex on a path in  $\mathcal{P}$  (Fig. 7a). To find  $u$  and  $w$ , our idea is to move over  $P^*$  starting from  $s$  and stop at the first such vertex  $u$  of  $P^*$ . Then use a modified DFS at  $u$  that runs over paths in  $\mathcal{P}$  only in reverse direction and over cross edges—but never over two subsequent cross edges<sup>1</sup>—and so explores subpaths of  $\mathcal{P}$  (marked orange in Fig. 7b). Whenever a vertex  $v$  of the clean area is reached, the DFS backtracks, i.e., the DFS assumes that  $v$  has no outgoing edges. We so guarantee that the clean area is not explored again and again. Afterwards we determine the latest vertex  $w$  on  $P^*$  that is a common vertex or has a cross edge with one of the explored subpaths. If such a vertex  $w$  is found, we either have a common subpath from  $u$  to  $w$  between a path  $P^c \in \mathcal{P}$  and  $P^*$  (which is removed as described above) or a deadlock cycle. If it is an extended deadlock cycle, we have a subpath on the cycle consisting of at least three vertices. As the proof of the following lemma shows the extended deadlock cycle can be destroyed by removing the inner vertices of the subpath and rerouting the paths via cross edges part of the extended deadlock cycle. We find an extended deadlock cycle by an additional run of the modified DFS from  $u$  to  $w$  (Fig. 7c) and reroute it along the cross edges. Figure 7d shows a rerouting where the path in  $R(\mathcal{P} \cup \{P^*\})$  starting with the vertices of the old path  $P_2$  becomes the “new” path  $P^*$ .

<sup>1</sup> This is due to a technical reason, the correctness paragraph in the proof of Lemma 11.

An  $s$ - $p$ -path  $\tilde{P}$  is called a *clean subpath* with respect to a set of good paths  $\mathcal{P}'$  if, for the extension  $P^{\text{ext}}$  of  $\tilde{P}$  by both an edge  $\{p, t\}$  and a vertex  $t$ ,  $\mathcal{P}' \cup \{P^{\text{ext}}\}$  is a set of good paths. The green path in Fig. 6 is a clean subpath.

The next lemma summarizes the computation of our rerouting  $R$ . We initially call the lemma with a clean subarea  $Q = \{s\}$  and a clean path that consists only of vertex  $p = s$ .

**Lemma 11** *There is an algorithm that accepts as input a valid path data scheme for  $\ell$  good  $s$ - $t$ -paths  $\mathcal{P}$ , a valid path data scheme for a possibly dirty  $s$ - $t$ -path  $P^*$ , as well as a clean area  $Q$  for  $\mathcal{P}$  including a clean  $s$ - $p$  subpath  $\tilde{P}$  of  $P^*$  and outputs a rerouting  $R$ , a clean area  $Q'$  including a clean  $s$ - $p'$  subpath  $\tilde{P}'$  with the properties that  $\tilde{P}$  is a subpath of  $\tilde{P}'$ , and  $|Q'| = |Q| + \Omega(n/\log k) \vee p' = t$ . The algorithm runs in  $O(nk^3 \log^3 k)$  time and uses  $O(n + k^2(\log k) \log n)$  bits.*

**Proof** We begin to describe our algorithm to detect common vertices and extended deadlock cycles. Afterwards we compute the rerouting  $R$  which is realized by  $(V^*, \vec{r}, \overleftarrow{r})$ . Initially,  $V^*$  is an array of bits consisting of all vertices of  $\mathcal{P} \cup \{P^*\}$ . During the algorithm we remove vertices from  $V^*$  whenever we change our paths. We want to use static space allocation to store the mappings  $\vec{r}$  and  $\overleftarrow{r}$ , which requires us to know the key set of the mapping in advance. We solve it by running the algorithm twice and compute the key set in the first run, reset  $Q$  and  $p$  from a backup, and compute the values and store them in a second run. For simplicity, we omit these details below and assume that we can simply store the values inside  $\vec{r}$  and  $\overleftarrow{r}$ .

Starting from vertex  $p$  we run along path  $P^*$ . We stop at the first vertex  $u$  if  $u$  is a common vertex of  $P^*$  and a path in  $\mathcal{P}$  (Fig. 7a) or it is an endpoint of a cross edge. For a simpler notion, we call such a vertex  $u$  of  $P^*$  a vertex *touching* the vertices in  $\mathcal{P}$ . Next we use a DFS without restorations (due to our details described below, the DFS stack consists of only  $O(\ell)$  vertices, which are incident to cross edges of the current DFS path). We so explore the vertices  $Q'$  from  $u$  that are reachable from  $u$  via edges on the paths  $\mathcal{P}$  used in reverse direction and cross edges, but ignore a cross edge if it immediately follows after another cross edge and ignore the vertices in the clean area  $Q$ . To allow an economical way of storing the stack, the DFS prefers reverse edges (edges on a path  $P \in \mathcal{P}$ ) compared to cross edges when iterating over the outgoing edges of a vertex. This guarantees that the vertices on the DFS-stack consist of at most  $\ell$  subpaths of paths in  $\mathcal{P}$ . We store the vertices  $Q'$  processed by the DFS in a choice dictionary [23, 25, 30] since a choice dictionary allows us to compute  $Q := Q \cup Q'$  in a time linear to the amount of elements in  $Q'$ . Moreover, we count the number  $q$  of vertices of  $P^*$  that touch  $Q'$ . Running over  $P^*$  starting from  $u$  we can determine the last such vertex  $w$  (Fig. 7b).

Then we run a modified DFS to construct a  $u$ - $w$ -path  $\tilde{P}$  that (1) consists only of vertices in  $Q'$ , that (2) uses no subsequent cross edges and that (3) uses every path in  $\mathcal{P}$  only once (Fig. 7c).

Note that these restrictions (1) – (3) still allow us to reach all vertices of  $Q'$ . To ensure restriction (3) we maintain a bit array  $F$  of *forbidden* paths where  $i \in F$  exactly if a subpath of  $P_i \in \mathcal{P}$  is part of the currently constructed path. To construct  $\tilde{P}$  the DFS starts at vertex  $u$  and processes a vertex  $v$  as follows:

1. Stop the DFS if  $w$  is reached.

2. // To guarantee restriction (1):  
 If  $v \notin Q'$ , then backtrack the DFS and color  $v$  white.
3. Iterate over all reversed edges  $(v, v')$ : recursively process  $v'$ .
4. // To guarantee restriction (2) and (3):  
 Iterate over all cross edges  $\{v, v'\}$ :  
 If  $\text{color}(v') \notin F$  and  $v'$  was not discovered by a cross edge, recursively process  $v'$ .
5. Backtrack the DFS to the predecessor  $\tilde{v}$  of  $v$  on  $\tilde{P}$ . If  $\{\tilde{v}, v\}$  is a cross-edge, color  $v$  white.  
 // Coloring  $v$  white guarantees that we can explore outgoing cross-edges of  $v$  if we reach  $v$  again over a reverse edges.

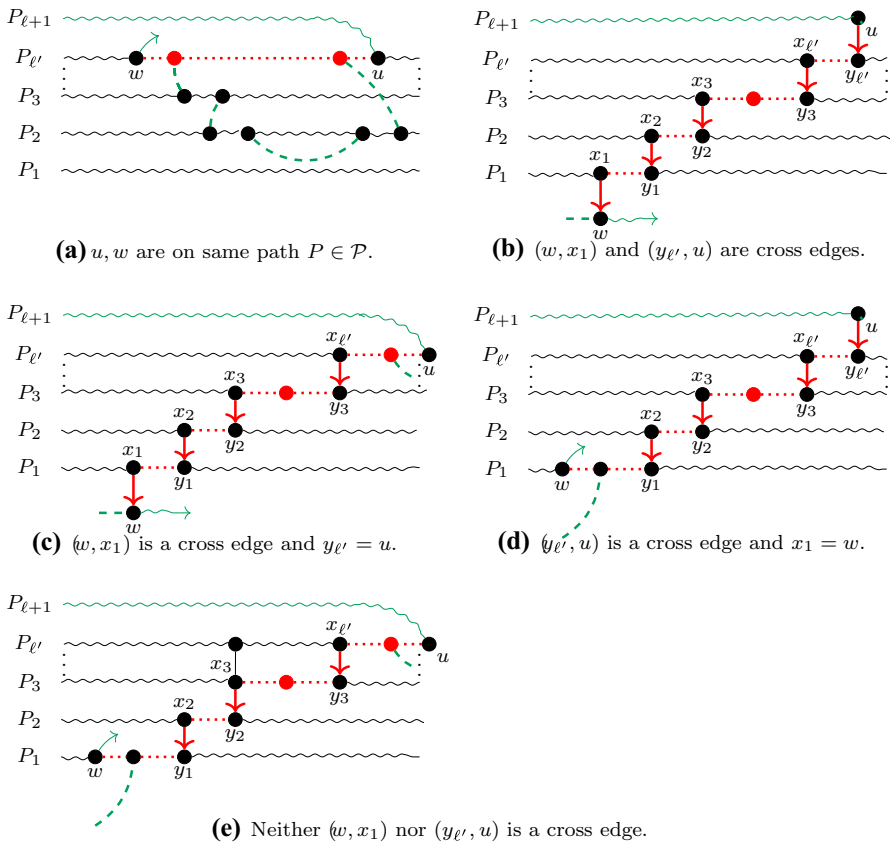
Let  $1, \dots, \ell$  be the colors of the paths in  $\mathcal{P}$  and  $(\ell + 1)$  be the color of  $P^*$ . In the case that  $u$  and  $w$  are both on one path of  $\mathcal{P}$ , we reroute the paths as follows. Set  $\vec{r}(u) = \overleftarrow{r}(w) = \text{color}(u)$ ,  $\overleftarrow{r}(u) = \vec{r}(w) = \ell + 1$  and update  $V^*$  accordingly, i.e., remove all internal vertices of the  $u$ - $w$  subpath of  $P^*$  from  $V^*$  (Fig. 8a). Afterwards, we set  $Q := Q \cup Q'$  and  $p = \text{next}(w)$  and repeat the whole algorithm above, unless  $\vec{r}$  is defined for  $\Theta(n/\log k)$  vertices ( $R$  uses  $O(n)$  bits) or we reached  $t$  while moving over  $P^*$ . In this case we are done and break the entire procedure.

Recall that the construction of the  $u$ - $w$ -path uses only one subpath of every path in  $\mathcal{P}$ . Without loss of generality, let  $\tilde{P}$  be  $w, x_1, \dots, y_1, x_2, \dots, y_2, x_3, \dots, y_{\ell'}, u$ , where  $(y_i, x_{i+1})$  are cross edges and  $x_i, \dots, y_i$  are subpaths of a path  $P_i \in \mathcal{P}$  (for some order of  $\mathcal{P}$  and some  $\ell'$  with  $0 < i < \ell' \leq \ell$ ).

We first assume that  $(w, x_1)$  and  $(y_{\ell'}, u)$  are both cross edges (Fig. 8b). If the  $x_i$ - $y_i$ -subpath of every  $P_i$  as well as the  $u$ - $w$ -subpath of  $P^*$  consists of exactly two vertices, then we found a simple deadlock cycle, in particular,  $u$  and  $w$  are the only vertices of  $P'$  that touch  $Q'$ . We keep the simple deadlock cycle by setting  $Q := Q \cup Q'$  and  $p = \text{next}(w)$  and repeat or break the whole algorithm analogously to the end of the previous paragraph.

Otherwise we have an extended deadlock cycle and we compute a rerouting as follows. For every cross edge  $(y_i, x_{i+1})$  with  $0 < i < \ell'$  on the DFS stack, we set  $\vec{r}(x_{i+1}) = i$ ,  $\overleftarrow{r}(y_i) = i + 1$ , remove all vertices of each path  $P_i$  in between  $x_i$  and  $y_i$  as well as all vertices of path  $P^*$  in between  $u$  and  $w$  from  $V^*$ . We now describe the rerouting at  $w$ . Vertex  $u$  is handled analogously. If as assumed  $(w, x_1)$  is a cross edge, set  $\vec{r}(x_1) = \ell + 1$  and  $\overleftarrow{r}(w) = 1$  (Fig. 8c). Otherwise,  $w$  is on path  $P_1$  before vertex  $y_1$ , or  $w = x_1$ . Then remove all vertices in between  $w$  and  $y_1$  on path  $P_1$  from  $V^*$  and set  $\vec{r}(w) = \ell + 1$  and  $\overleftarrow{r}(w) = 1$  (Fig. 8d). For the case where neither  $(w, x_1)$  nor  $(y_{\ell'}, u)$  is a cross edge, as well as  $u$  and  $w$  are on different paths, see Fig. 8e for the rerouting. For all cases described in this paragraph, finally set  $Q := Q \cup Q'$  and repeat or break the algorithm exactly as described at the end of the previous paragraph.

Note that in all cases above, we only add information for  $O(k)$  predecessors and successors with each extension of  $Q$  to our rerouting. Thus, our rerouting information increases in small pieces. Whenever  $\vec{r}$  is defined for  $\Theta(n/\log k)$  vertices ( $R$  uses  $O(n)$  bits) or we reached  $t$  while moving over  $P^*$ , we break the algorithm above. By applying Lemma 5 on each rerouted path (the graph induced by the vertices of the paths is taken as input graph), we can make all paths chordless without introducing further



**Fig. 8** Sketch of rerouted paths to remove common vertices and extended deadlock cycles. The dotted and dashed subpaths are removed from the paths and the red lines with an arrow show a switch to another path (Color figure online)

cross edges or extended deadlock cycles. Finally, return the rerouting  $(V^*, \vec{r}, \overleftarrow{r})$  as well as the new clean area  $Q$  and the end vertex  $p$  of a new clean path in  $R(\mathcal{P} \cup \mathcal{P}')$ .

**Correctness** One can see that our algorithm discovers every common subpath of  $P^*$  with a path  $P^c \in \mathcal{P}$  as well as every extended deadlock cycle since, in both cases, we have a vertex  $u$  followed by  $w$  on  $P^*$  such that  $w$  is discovered from  $u$  while running backwards on  $\mathcal{P}$ . Since  $V^*$  only shrinks (no new cross edges can occur), since by construction the new paths have no common vertices accept  $s$  and  $t$  (here it is important that the constructed  $u$ - $w$ -path has no subsequent cross edges) and since we choose  $w$  as the latest vertex of  $P^*$  touching  $Q'$ ,  $Q$  maintains a clean area when adding  $Q'$  to it. Note that each vertex with a rerouting information becomes part of  $Q$ . Thus, the size of  $Q$  increases by  $\Omega(n/\log k)$  when our algorithm stops unless we reach  $t$  while running on  $P^*$ .

**Efficiency** Concerning the running time, note that we travel along  $P^*$  as well as construct a set  $Q'$  with a DFS on the paths in  $\mathcal{P}$ . We process all vertices of  $Q'$  at most



once by each of the two standard DFS runs (once to construct  $Q'$  as well as once to find a  $u$ - $w$ -path for computing the rerouting) before  $Q'$  is added into  $Q$ . Within the processing of each vertex  $v$ , we call  $\text{prev}(v)$  only once.

Since each access to a vertex  $v$  can be realized in  $O(\text{deg}(v) + k^3 \log^3 k)$  time (Lemma 7), the total running time is  $O(nk^3 \log^3 k)$ .

Concerning the space consumption observe that  $\vec{r}$  and  $\overleftarrow{r}$ , a backup of  $Q$  and  $p$  as well as all arrays use  $O(n)$  bits. To run the DFS,  $O(n + \ell \log n)$  bits suffice (storing the colors white, grey and black in  $O(n)$  bits and  $O(\ell \log n)$  bits for the DFS stack since we only have to store the endpoints of  $O(\ell)$  subpaths that are part of the stack). Together with the space used by Lemma 7, our algorithm can run with  $O(n + k^2(\log k) \log n)$  bits. □


We can not store our rerouted paths directly in a valid path data scheme, with the reasons outlined shortly. Instead, to follow the path with respect to the rerouting we provide a so-called weak path data structure for all  $s$ - $t$ -paths of  $R(\mathcal{P} \cup \{P'\})$ , defined in the following.

Recall that we have a path data scheme storing the single path  $P^*$  and a path data scheme storing the set  $\mathcal{P}$  of paths. Let  $Q, p$  as well as  $R$ —realized by  $(V^*, \vec{r}, \overleftarrow{r})$ —be returned by the previous lemma. Our valid path data scheme can only store  $s$ - $t$ -paths, but  $R(\mathcal{P} \cup \{P'\})$  may contain an  $s$ - $p'$ -path  $P''$  with  $p' \neq t$ . In the case that  $p' \neq t$ , we compute a path numbering and a valid path data scheme for the path  $P^{**}$  obtained from  $P^*$  by replacing the subpath  $P'$  by  $P''$ . (The path  $P^{**}$  replaces the path  $P^*$  in the computation of the next batch.) Furthermore, we remove the path  $P''$  from  $R(\mathcal{P} \cup \{P'\})$  by removing its vertices from  $V^*$ . It is easy to see that this modification of the paths can be done in the same bounds as stated in Lemma 11.

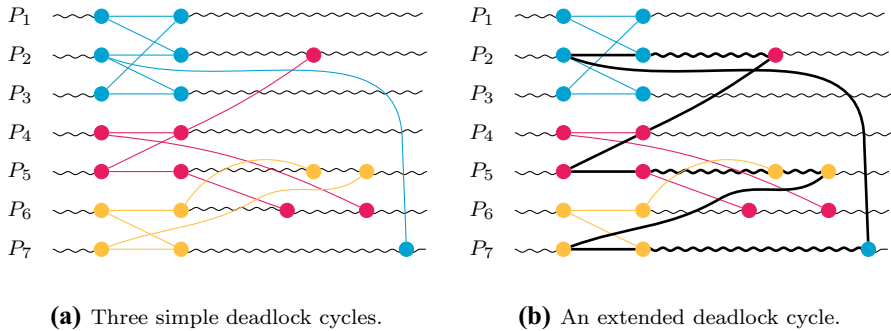
After the modification of the paths we show how to follow the paths with respect to  $R$ . Since we can find out in  $O(1)$  time if a vertex  $u$  belongs to  $P^*$ , a call of  $\text{prev}$  and  $\text{next}$  below is forwarded to the correct path data scheme. We now overload the  $\text{prev}$  and  $\text{next}$  operations of our path data structure as follows, and so get a *weak path data structure* for  $R(\mathcal{P} \cup \{P'\}) \setminus \{P''\}$  that supports  $\text{prev}'$  and  $\text{next}'$  in time  $O(\text{deg}(v) + k^3 \log^3 k)$  by using  $O(n + k^2(\log k) \log n)$  bits and that supports access to  $V^*$  in constant time. The runtime is due to the fact that we realize  $\text{prev}'$  and  $\text{next}'$  by computing the paths within at most one region, from which we know the color of all neighbors of  $u$ . In detail, the weak path data structure supports the following operations where  $u \in V^*$ .

- $\text{prev}'(u)$ : If  $\overleftarrow{r}(u) = 0$ , return  $\text{prev}(u)$ . Else, return  $v \in N(u) \cap V^*$  with  $\text{color}(v) = \overleftarrow{r}(u)$ .
- $\text{next}'(u)$ : If  $\vec{r}(u) = 0$ , return  $\text{next}(u)$ . Else, return  $v \in N(u) \cap V^*$  with  $\text{color}(v) = \vec{r}(u)$ .
- $\text{invStar}(u)$ : Returns true exactly if  $u \in V^*$ .

Note that the key difference between a weak path data scheme and a valid path data scheme is that we are unable to provide a  $\text{color}$  operator without fully traversing all paths, which of course takes too long. To compute a valid path data scheme for  $\ell \in \mathbb{N}$  good  $s$ - $t$ -paths we have to partition  $G[V^*]$  into regions via set a boundary vertices  $B' \subseteq V^*$ . To construct  $B'$  we start with the set  $S$  of neighbors of  $s$  that belong to  $V^*$  and



**Fig. 9** Our algorithm determines a boundary by moving from  $s$  along all paths with respect to the cross edges (colored edges), i.e., it cannot move from a vertex that is an endpoint of a cross edge if the other endpoint is not yet reached. In the example, we assume that  $k \log k = 18$ . The number inside the vertices denotes the seen vertices since adding vertices to the boundary the last time (Color figure online)



**Fig. 10** Illustration of the construction of an extended deadlock cycle (Color figure online)

move  $S$  to the right “towards”  $t$  such that  $S$  always is a separator that disconnects  $s$  and  $t$  in  $G[V^*]$ . Whenever there are  $O(k \log k)$  vertices “behind” the previous separator that was added to our boundary, we add  $S$  to our boundary—see also Fig. 9. While moving  $S$  towards  $t$ , we always make sure that the endpoints of a cross edge are not in different components in  $G[V^* \setminus S]$ . Again, this can be thought of as a sliding window algorithm. This means, we do not move  $S$  over an endpoint of a cross edge unless the other endpoint is also in  $S$ . We show in the next lemma that this is possible if we have no extended deadlock cycles with the help of the following structural lemma, for which we define the following notion. Given a simple deadlock cycle  $Z$  consider a partition of the vertices of  $Z$  in two sets  $Z_{\text{left}}$  and  $Z_{\text{right}}$  such that each  $v \in Z_{\text{left}}$  has a successor in  $Z_{\text{right}}$ . Such a partition is possible for every simple deadlock cycle.

**Lemma 12** *Assume that we have no extended deadlock cycles in  $G$  with respect to  $\mathcal{P}$ . Then there is a partial ordering among the simple deadlock cycles defined as follows: A simple deadlock cycle  $Z$  is smaller than a simple deadlock cycle  $Z'$  if there is a path  $P \in \mathcal{P}$  with vertices  $u$  before  $v$  such that (1)  $u$  in  $Z_{\text{left}}$  and (2) there is a cross edge  $\{v', v\}$  with  $v' \in Z'_{\text{left}}$ .*

**Proof** For a cross edge  $e = \{v', v\}$  with the property of the lemma we say that  $e$  is a cross edge of  $Z'$  that jumps behind  $Z$ . Assume that the ordering is not well defined. Then there is a sequence of  $x \geq 1$  simple deadlock cycles  $Z_1, Z_2, \dots, Z_x$  such that  $Z_1$  jumps behind  $Z_x$  and, for  $1 \leq i < x$ ,  $Z_{i+1}$  has a cross edge that jumps behind  $Z_i$ . Figure 10a shows an example for  $x = 3$  with the two different kinds of interaction between two simple deadlock cycles: a cross edge part of the simple deadlock cycles

is used to jump on a common paths  $P \in \mathcal{P}$  (see the red and orange simple deadlock cycles) or another *extra cross edge* is used (see the green and red simple deadlock cycles). It is easy to see that we then can use this sequence of simple deadlock cycles and possibly some extra cross edges to build one or more extended deadlock cycles; a contradiction.

We use the following strategy for resolving simple deadlock cycles: move  $S$  one vertex further along the path for all vertices part of the simple deadlock cycle.

In the case that the previous lemma returned paths containing an  $s$ - $p'$ -path with  $p' \neq t$ , then our weak path data structure gives access to  $\ell$  good  $s$ - $t$ -paths. Otherwise,  $p' = t$  and we have access to  $\ell := \ell + 1$  good  $s$ - $t$ -paths.

**Lemma 13** *Given a weak path data structure for accessing  $\ell$  good  $s$ - $t$ -paths  $\mathcal{P}'$  there is an algorithm that computes a valid path data scheme storing  $\ell$  good  $s$ - $t$ -paths. The algorithm runs in  $O(nk^3 \log^3 k)$  time and uses  $O(n + k^2(\log k) \log n)$  bits.*

**Proof** Let  $V^*$  be the vertices of  $P^*$  and  $\mathcal{P}'$  obtained from the weak path data structure. As in the last proof, we assume that the paths in  $\mathcal{P}'$  are colored from  $1, \dots, \ell$ . For the case that  $p' \neq t$ , we compute a new path data scheme for  $P^*$  by Lemma 8. Moreover, we remove the vertices of  $P^*$  from  $V^*$  and start to compute our path data scheme for  $\mathcal{P}'$ .

Our algorithm works in three steps. First we compute the boundary, then a path data scheme, and finally make it valid.

**Compute the boundary** Start with the set  $S$  of all  $|\mathcal{P}'|$  neighbors  $u$  of  $s$  in  $G[V^*]$ . Note that the vertices in  $S$  are candidates for a boundary because our paths are chordless and thus  $S$  cuts all  $s$ - $t$ -paths; in other words, there is no path in  $G[V^* \setminus S]$  that connects  $s$  and  $t$ . For performance reasons, we cache  $\text{next}'(u)$  for each vertex  $u$  in  $S$  as long as  $u$  is in  $S$  so that we have to compute it only once. From  $S$  we can compute the next set  $S'$  of candidates for a boundary. The idea is to try to put for each vertex  $u \in S$  the vertex  $v = \text{next}'(u)$  of the same path into  $S'$ . However, we cannot simply do this if  $u$  has a cross edge that can bypass  $S'$ . In that case we take  $u$  into  $S'$ . If  $S = S'$ , then one (or multiple) deadlock cycles restrain us to add new vertices to  $S'$ . We are either able to take the next vertex of every vertex of  $S$  that is part of a simple deadlock cycle into  $S'$  or show that an extended deadlock cycle exists, which is a contradiction to our property that our paths are good.

We now focus on the structures that we use in the subsequent algorithm. Let  $D$  (initially  $D := S \cup \{s\}$ ) be a set consisting of all vertices of the connected component of  $G[V^* \setminus S]$  that contains  $s$  as well as the vertices of  $S$ . (Intuitively speaking,  $D$  is the set of vertices for which the membership in a boundary was already decided.) We associate a vertex with the index of the path to which it belongs. Technically we realize  $S$  and  $S'$  as arrays of  $|\mathcal{P}'|$  fields such that  $S[i]$  and  $S'[i]$  store vertices  $u_i$  and  $v_i$ , respectively, of  $P_i \in \mathcal{P}'$  ( $1 \leq i \leq |\mathcal{P}'|$ ). Moreover, for each  $u_i \in S$  on path  $P_i \in \mathcal{P}'$  we store the number  $\text{d}(i) = |(N(u_i) \cap V^*) \setminus D|$  in an array of  $|\mathcal{P}'|$  fields. These numbers represent the number of edges whose head points at a vertex  $u_i$  of  $P_i$  outside of  $D$ . If  $\text{d}(i) \neq 1$ , then  $u_i$  has a cross edge to some vertex not in  $D$  and we cannot simply put  $v_i := \text{next}'(u_i)$  inside  $S'$ , because we might get a cross edge that

can be used by an  $s$ - $t$ -path in  $G[V^* \setminus S']$ . Otherwise, if  $d(i) = 1$ , then  $u_i$  has no cross edge possibly bypassing  $S'$  and we can take  $v_i$  into  $S'$ . Let  $B'$  be an initially empty set of boundary vertices. To know when the next set  $S$  can be added to the boundary  $B'$  while ensuring that  $B'$  defines regions of size  $O(k \log k)$ , we use a counter  $c$  (initially  $c := 0$ ) to count the number of vertices that we have seen until extending the boundary the last time.

We now show our algorithm to compute  $S'$  given  $S$ . Take  $B' := S$ , initialize  $d(i)$  as  $|(N(u_i) \cap V^*) \setminus D|$  for each  $u_i \in S$ , and compute a set  $S'$  in phases as follows. For each  $u_i \in S$ , if  $d(i) = 1$ , take  $v_i = \text{next}'(u_i)$  into  $S'$ , increment  $c$  by one, and set  $d(i) = |(N(v_i) \cap V^*) \setminus D|$ . Moreover, for each  $w_j \in (N(v_i) \cap S)$  decrement  $d(j)$  by one. If  $d(i) \neq 1$ , take  $u$  into  $S'$ . If  $S \neq S'$ , add  $S'$  into  $D$ , set  $S := S'$  and repeat the phase.

Otherwise (i.e.,  $S = S'$ ) we run into the *special case* where we have encountered a deadlock cycle (possible a combination of deadlock cycles), which we can handle as follows. Since  $\mathcal{P}'$  is good and  $\tilde{P}$  is clean, we have one or more simple deadlock cycles with vertices in  $S$ . Set  $c = c + |\mathcal{P}'|$  and take  $\tilde{S} = \{\text{next}'(u) \mid u \in S\}$ . Note that all these simple deadlock cycles have only vertices in  $S \cup \tilde{S}$ , but not all vertices of  $\tilde{S}$  have to be part of a simple deadlock cycle. To compute  $S'$  from  $\tilde{S}$ , we have to remove the vertices from  $\tilde{S}$  that are not part of the simple deadlock cycle, which we can do as follows. Take  $S' := \tilde{S}$ . Check in rounds if  $S'$  separates  $S$  from  $t$ , i.e., for each  $v_i \in S'$  with  $1 \leq i \leq |\mathcal{P}'|$ , choose  $u_i \in S$  and check if  $d(j) > \ell$ . If not, additionally check if a  $w_j \in (N(u_i) \cap V^*) \setminus (D \cup \{S'[j]\})$  exists. If one of the check passes, some edge of  $u_i$  bypasses  $S'$  and we replace  $v_i$  by  $u_i$  in  $S'$  and decrement  $c$  by one. We then say that we *withdraw* from  $v_i$ . Then repeat the round until no such  $w$  can be found for any vertex. Afterwards, set  $S'$  as the new  $S$ , add  $S$  into  $D$ , and compute the  $d$ -values for the new vertices in  $S$ , and proceed with the next phase.

If  $c > k \log k$  after adding  $S$  into  $D$ , add the vertices of  $S$  to  $B'$ , and reset  $c := 0$ . Repeat the algorithm until  $S = N(t) \cap V^*$  and add the vertices of  $S$  to  $B'$ . Finally, add all vertices of  $V^*$  with large degree to  $B'$ .

**Correctness of the boundary** We show the following invariant whenever we update  $S$ : the removal of  $S$  in  $G[V^*]$  disconnects  $s$  and  $t$ . Initially, this is true for the neighbors of  $s$  because their removal disconnects each chordless path in  $G[V^*]$ . If we have a vertex  $u \in S$  on  $P_i \in \mathcal{P}$  with only one neighbor in  $G[V^* \setminus D]$ , i.e., we have  $d(i) = 1$ , then  $S' = (S \setminus \{u\}) \cup \{\text{next}'(u)\}$  is again a valid separator.

If we have no such vertex  $u$ , then we construct a set  $S'$  and we check for each vertex in  $S$  if it has a cross edge that destroys  $S'$  being a separator, i.e., we have an  $s$ - $t$  path using the cross edge, but no vertex of  $S'$ . We only set  $S := S'$  if  $S'$  is again a separator.

We next show that we make progress with each separator  $S$  over all phases, i.e., we find a new separator so that  $D$  increases. As long as we do not run into the special case having a deadlock cycle, we can clearly make progress. Let  $S'$  be defined as in the beginning of the special case, i.e.,  $S'$  consists of the successors on the paths of each vertex in  $S$ . In the following we show that we only enter the special case exactly if there exists a simple deadlock cycle  $Z$  with (Property 1)  $Z_{\text{left}} \subseteq S$  and (Property 2) there exists no cross edge  $\{v_1, v_2\}$  with  $v_1 \in Z_{\text{left}}$  and  $v_2 \notin Z_{\text{right}}$ . We call a simple deadlock cycle for which these properties hold a *special deadlock cycle*.

Since we are in the special case we know that all vertices of  $S$  are part of some deadlock cycle. By Lemma 12 we have a partial ordering  $\phi$  among all deadlock cycles having a vertex in  $S$ . Then for a smallest deadlock cycle  $Z$  with respect to  $\phi$  we must have  $Z_{\text{left}} \subseteq S$ , i.e., Property 1 of a special deadlock cycle holds for  $Z$ . Since  $Z$  is a smallest deadlock cycle with respect to  $\phi$  there can not be a cross edge described in Property 2 of a special deadlock cycle. Now, for each  $u \in Z_{\text{left}}$ ,  $v$  being a successor of  $u$  and for each cross edge  $\{u, u'\}$ ,  $v$  and  $u'$  are in  $Z_{\text{right}} \subseteq S'$  and we do not have to withdraw any vertex of  $Z_{\text{right}}$  from  $S'$ , and thus we make progress also in the special case.

**Computing the new valid path data scheme** We already know  $V^*$  and the new boundary  $B'$ . For a valid path data scheme it remains to compute a path numbering  $A'$ , the colors for the boundary vertices and their predecessors and successors that we store in an  $O(n)$ -bit structure  $C'$  using static space allocation. We can compute  $A'$  by moving along each path from  $s$  to  $t$ .

It is important that all vertices of an  $s$ - $t$ -path are colored with the same color, however we want to use Lemma 6 that computes path inside one region. We first give an intuition of our approach and describe the details in the subsequent paragraphs. Since  $s$ - $t$ -paths are disconnected by boundary vertices our approach is to start with an arbitrary coloring of the neighbors  $W$  of  $s$  part of the paths that are the first boundary vertices and from these make a parallel run over the paths. Whenever we enter a region we explore it with our fixed deterministic network-flow algorithm and compute paths within the region. Using the paths we propagate the colors of the boundary vertices along the paths to the next boundary vertices, their predecessors and successors and repeat our algorithm.

Let  $D \subseteq V^*$  (initially  $D := W \cup \{s\}$ ) be the set of vertices for that we have already decided their color. To find a fixed routing within a region we explore the region of every neighbor of  $W$  in  $G[V^* \setminus (B' \cup D)]$  using a BFS that skips over vertices in  $B' \cup D$ . We collect all visited vertices of one region in a set  $U$  realized by a balanced heap. (If  $U$  is empty, the region was already explored from another neighbor of a vertex of  $W$  and we continue with the next vertex  $v \in W$ .) In addition we construct two sets  $S'$  and  $T'$ . Each visited vertex  $u \in U$  is put in  $S'$  if  $\text{prev}'(u) \in B'$  holds and put in  $T'$  if  $\text{next}'(u) \in B'$  holds (a vertex  $u$  can be part of  $S'$  as well as  $T'$  if both holds). We apply Lemma 6 with  $U$ ,  $S'$  and  $T'$  as input and get paths connecting each vertex of  $S'$  with another vertex of  $T'$ . We set  $D := D \cup U$  to avoid the exploration of explored regions again.

Let  $F$  be a set of vertices whose successors of  $W$  are in  $U$ . Intuitively, by having  $F$  we can avoid computing the paths of a region again and gain. For each vertex  $v \in F$  we use  $\text{next}'(v)$  and so move to a vertex  $u$  of  $S'$ —we are now on a path computed by Lemma 6. We determine  $q$  as the color of  $v$ , store  $q$  as color for  $u$ , run along the path until it ends. We store the color  $q$  for the last vertex  $w$  of the path and use  $\text{next}'(w)$  to move to a boundary vertex  $v'$  that we also color with  $q$ . If the vertex  $u' = \text{next}'(v')$  is on a computed path we repeat this paragraph with  $v := v'$ . Otherwise, we put  $u'$  into a set  $W'$  and remove  $v$  from  $F$ . If  $F$  loses its last vertex we set  $W := W'$  and  $F = W' = \emptyset$  and repeat the iteration. Otherwise, we proceed with the next vertex  $v \in W$ .

While moving over the paths we extend  $D$  by the visited vertices. Moreover we compute the path numbering for  $A'$  and all colors are stored in  $C'$ . Recall that in order to use static space allocation we first need to know the set of all vertices in advance that we afterwards want to color. Do get that set we can run the algorithm in two steps. In the first we do not store the colors but only compute a set of vertices that must be colored. Using it as a key set for static space allocation we repeat our algorithm in a second step and store the coloring.

The triple  $(A', B', C')$  represents the new valid path data structure.

**Correctness of the new valid path data scheme** The boundary and the rerouted paths represented by the weak path data structure fixes the boundary vertices and their predecessor and successors. Hence our selection of the vertices for  $U, S'$  and  $T'$  is also fixed for each region. Computing the paths using Lemma 6 fixes the path numbering. Since a (re)computation with the same algorithm produces the same paths (being subpaths of our  $s$ - $t$ -paths) and since we used the paths to propagate the coloring, the endpoints of a path and thus, all all colored vertices along a path are equally colored in  $C'$ . Therefore our path data scheme is valid. Moreover,  $S'$  and  $T'$  consists of the endpoints of subpaths of good paths. Hence, a network-flow algorithm can connect each vertex of  $S'$  with another vertex of  $T'$  and by Lemma 4 our paths are chordless and deadlock free paths.

**Efficiency** We first consider the time to compute the boundary. For each new vertex  $u$  in  $S$ , we spend  $O(\text{deg}(u))$  time to determine once its degree in  $G[V^* \setminus D]$  and to decrement already stored degrees of the neighbors of  $u$ . In each phase we iterate over all elements of  $S$  and make  $|S\Delta S'| = \Omega(1)$  progress on at least one path, which means that we need at most  $O(n)$  phases. Determining the previous / next vertex  $u$  on a path can be done in time  $O(\text{deg}(u) + k^3 \log^3 k)$ .

Dealing with simple deadlock cycles requires us to compute the  $\ell$  initial members of  $S'$ , which cost us  $O(\ell(\text{deg}(u) + k^3 \log^3 k))$  time and roundwise replace members in  $S'$  by their predecessors. Because  $|S'| = O(\ell)$ , this can cause  $O(\ell)$  rounds. We spend  $O(\ell^2)$  time to check if  $S'$  is a separator. Thus,  $O(\ell^3)$  time allows us to compute  $S'$  in the special case, and  $O(\ell)$  time allows us to add the at most  $\ell$  vertices to  $D$  within each phase.

To construct a valid path data scheme we need to explore all regions once, which can be done in time linear to the size of each region plus the number of edges that leave the region, i.e., in  $O(nk)$  total time. The construction of  $S'$  and  $T'$  requires the execution of  $\text{prev}'$  and  $\text{next}'$  on each vertex once, which runs in  $O(nk^3 \log^3 k)$  time. Lemma 6 uses  $O(n'k^2 \log^2 k)$  time for a region of size  $n'$ . Using it on every region once can be done in  $O(nk^2 \log^2 k)$  time. Propagating the color requires us to move along the paths, which can be done in total time  $O(nk^3 \log^3 k)$  time. Thus, our algorithm runs in time  $O(nk^3 \log^3 k)$ .

Our space bound is mainly determined by  $D, V^*, A', B'$  and  $C'$  each of  $\Theta(n)$  bits and the application of Lemma 7 ( $O(k^2(\log k) \log n)$  bits), which is used in the weak path data structure whenever we access a previous or next vertex on a path. For  $S, W$  and  $F$  (and their copies) we need  $\Theta(\ell \log n)$  bits. The construction of the paths inside regions uses  $O(k^2(\log k) \log n)$  bits (Lemma 6). In total, we use  $O(n + k^2(\log k) \log n)$  bits. □

Given path storage scheme storing  $\ell = O(k)$  good  $s$ - $t$ -paths and an  $(\ell + 1)$ th path  $P^*$  we can batchwise compute a path storage scheme storing  $(\ell + 1)$  good  $s$ - $t$ -paths by subsequently executing  $O(\log k)$  times Lemma 11 and Lemma 13.

**Corollary 1** *Given a path storage scheme storing  $\ell = O(k)$  good  $s$ - $t$ -paths  $\mathcal{P}$  and an  $(\ell + 1)$ th dirty path with respect to  $\mathcal{P}$  we can compute a path storage scheme storing  $(\ell + 1)$  good  $s$ - $t$ -paths in  $O(nk^3 \log^3 k)$  time using  $O(n + k^2(\log k) \log n)$  bits.*

Initially our set of  $s$ - $t$ -paths  $\mathcal{P}$  is empty. Then a repeated execution of Corollary 1 allows us to show our final theorem, which will be used to find separators of size  $O(k)$ . Note that we gain an extra factor of  $(k \log k + \log^* n)$  in the runtime due to non-constant time access to the graph interface presented to the DFS, with the  $\log^* n$  factor coming from the use of the space-efficient DFS of Lemma 1.

**Theorem 3** *Given an  $n$ -vertex graph  $G = (V, E)$  with treewidth  $k$  and two vertices  $s, t \in V$  there is an algorithm that can compute a maximum amount of  $O(k)$  chordless vertex-disjoint paths from  $s$  to  $t$  in  $O(n(k \log k + \log^* n)k^4 \log^3 k)$  time using  $O(n + k^2(\log k) \log n)$  bits.*

Knowing a maximum set of  $\ell = O(k)$  vertex-disjoint paths between two vertices  $s$  and  $t$ , we can easily construct a vertex separator for  $s$  and  $t$ .

**Corollary 2** *Given an  $n$ -vertex graph  $G = (V, E)$  with treewidth  $k$ , and two vertices  $s \in V$  and  $t \in V$ ,  $O(n(k \log k + \log^* n)k^4 \log^3 k)$  time and  $O(n + k^2(\log k) \log n)$  bits suffice to construct a bit array  $S$  marking all vertices of a vertex separator of size  $O(k)$  for  $s$  and  $t$ .*

**Proof** First construct the maximum number  $O(k)$  of possible chordless pairwise vertex-disjoint paths from  $s$  to  $t$  (Theorem 3). Then try to construct a further  $s$ - $t$ -path with a DFS as describe in the proof of Lemma 10 and store the set  $Z$  of vertices that are processed by the DFS. Finally run along the paths from  $s$  to  $t$  and compute the set  $S$  consisting of the last vertex on each path that is part of  $Z$ . These vertices are an  $s$ - $t$ -separator.  $\square$

Many practical applications that use treewidth algorithms have graphs with treewidth  $k = O(n^{1/2-\epsilon})$  for an arbitrary  $\epsilon > 0$ , and then our space consumption is  $O(n)$  bits.

## 4 Sketch of Reed's Algorithm

In this section we first sketch Reed's algorithm to compute a tree decomposition and then the computation of a so-called balanced  $X$ -separator. In the following sections, we modify his algorithm to make it space efficient.

Reed's algorithm [37] takes an undirected  $n$ -vertex  $m$ -edge graph  $G = (V, E)$  with treewidth  $k$  and an initially empty vertex set  $X$  as input and outputs a balanced tree decomposition of width  $8k + 6$ . To decompose the tree he makes use of separators. An  $X$ -separator is a set  $S \subset V$  such that  $S$  separates  $X$  among the connected components

of  $G[V \setminus S]$  and no component contains more than  $2/3|X|$  vertices of  $X$ . A balanced  $X$ -separator  $S$  is an  $X$ -separator with the additional property that no component of  $G[V \setminus S]$  contains more than  $2/3|V|$  vertices.

The decomposition works as follows. If  $n \leq 8k + 6$ , we return a tree decomposition  $(T, B)$  consisting of a tree with one node  $r$  (the root node) and a mapping  $B$  with  $B(r) = V$ . Otherwise, we search for a so-called balanced  $X$ -separator  $S$  of size  $2k + 2$  that divides  $G$  such that  $G[V \setminus S]$  consists of  $x \geq 2$  connected components  $\Gamma = \{G_1, \dots, G_x\}$ . Then, we create a new tree  $T$  with a root node  $r$ , a mapping  $B$ , and set  $B(r)$  to  $X \cup S$ . For each graph  $G_i \in \Gamma$  with  $1 \leq i \leq x$ , we proceed recursively with  $G' = G_i[V(G_i) \cup S]$  and  $X' = ((X \cap V(G')) \cup S)$ . Every recursive call returns a tree decomposition  $(T_i, B_i)$  ( $i = 1, \dots, x$ ). We connect the root of  $T_i$  to  $r$ , we then set  $B(w) = B_i(w)$  for all nodes  $w \in T_i$ . After processing all elements of  $\Gamma$  return the tree decomposition  $(T, B)$ .

Since a balanced  $X$ -separator is used, the tree has a depth of  $O(\log n)$ , and thus the recursive algorithm produces at most  $O(\log n)$  stack frames on the call stack—each stack frame is associated with a node  $w$  of  $T$ . A standard implementation of the algorithm needs a new graph structure for each recursive call. In the worst-case, each of these graphs contains  $2/3$  of the vertices of the previous graph. Thus, the graphs on the stack frame use  $\Theta((n + m) \log n) = \Theta(kn \log n)$  bits. Storing the tree decomposition  $(T, B)$  requires  $\Theta(kn \log n)$  bits as well. The various other structures needed can be realized within the same space bound. In conclusion, a standard implementation of Reed’s algorithm requires  $\Theta(kn \log n)$  bits.

The next lemma shows a space-efficient implementation for finding a balanced  $X$ -separator.

**Lemma 14** *Given an  $n$ -vertex graph  $G = (V, E)$  with treewidth  $k$  and  $X \subseteq V$  of at most  $6k + 6$  vertices, there is an algorithm for finding a balanced  $X$ -separator of size  $2k + 2$  in  $G$  that, for some constant  $c$ , runs in  $O(c^k n \log^* n)$  time and uses  $O(n + k^2(\log k) \log n)$  bits.*

**Proof** We now sketch Reed’s ideas to compute a balanced  $X$ -separator. To compute a balanced  $X$ -separator we compute first an  $X$ -separator  $S_1$ . To make it balanced, we compute an additional  $R$ -separator  $S_2$  where  $R$  is a set of vertices that is in some sense equally distributed in  $G$ . Then  $S = S_1 \cup S_2$  is a balanced  $X$ -separator.

Reed computes an  $X$ -separator by iterating over all  $3^{|X|}$  possibilities to split  $X$  into three vertex-disjoint sets  $X_1, X_2 \subseteq V$  and  $X_S$  with  $|X_S| \leq k$  and  $|X_1|, |X_2| \leq \max\{k, 2/3|X|\}$ .

For each iteration connect two new vertices  $s'$  and  $t'$  with all vertices of  $X_1$  and  $X_2$ , respectively, compute vertex disjoint paths with Corollary 2 to find a separator  $S$  and check if  $X_S \subseteq S$  holds.

We now shortly describe Reed’s computation of the set  $R$ . Run a DFS on the graph  $G$  and compute in a bottom up process for each vertex  $v$  of the resulting DFS tree the number of descendants of  $v$ . Whenever this number exceeds  $n/(8k + 6)$ , add  $v$  to the initially empty set  $R$  and reset the number of descendants of  $v$  to zero. At the end of the DFS, the set  $R$  consist of at most  $8k + 6$  vertices, which can be used to compute a balanced  $X$ -separator as described above. Similar to the  $X$ -separator, we compute an  $R$ -separator.



We now show how to compute  $R$  using  $O(n)$  bits. To prove the lemma, we use the following observation. Whenever the number of descendants for a node  $u$  is computed, the numbers of  $u$ 's children are not required anymore.

The idea is to use a *balanced parentheses* representation, which consists of an open parenthesis for every node  $v$  of a tree, followed by the balanced parentheses representation of the subtrees of every child of  $v$ , and a matching closed parenthesis.

Consequently, if  $v$  is a vertex with  $x$  descendants having its open parenthesis at position  $i$  and its closed parenthesis at position  $j$ , then the difference between  $i$  and  $j$  is  $2x$ .

Note that, taking an array  $A$  of  $2n$  bits, we can store the number of descendants of  $v$  in  $A[i \dots j]$  as a so-called self-delimiting number by writing  $x$  as  $1^x0$ . This means that we overwrite the self-delimiting numbers stored for the descendants of the children of  $v$ .

To construct  $R$  we run a space-efficient DFS twice, first to construct a balanced parentheses representation of the DFS tree, which is used to compute the descendants of each vertex in the DFS tree and so choose vertices for the set  $R$ , and a second time to translate the IDs of the chosen vertices since the balanced parentheses representation is an ordinal tree, i.e., we lose our original vertex IDs and the vertices get a numbering in the order the DFS visited the vertices. After choosing the vertices that belong to the set  $R$  and marking them in a bit array  $R'$ , we run the DFS again and create a bit array  $R^*$  that marks every vertex  $v$  that the DFS visits as the  $i$ th vertex if and only if  $i$  is marked in  $R'$ .

It remains to show how to compute the bit array  $R'$ . Let  $P$  be a bit array of  $2n$  bits storing the balanced parentheses representation, and let  $A$  be a bit array of  $2n$  bits that we use to store the numbers of descendants for some vertices. Note that a leaf is identified by an open parenthesis followed by an immediately closed parenthesis. Moreover, since the balanced representation is computed via a DFS in pre-order, we will visit the vertices by running through  $P$  in the same order. Note that Munro and Raman [36] showed a succinct data structure for balanced representation that initializes in  $O(n)$  time and allows to compute the position of a matching parenthesis, i.e., given an index  $i$  of an open (closed) parenthesis there is an operation  $\text{findclose}(i)$  ( $\text{findopen}(i)$ ) that returns the position  $j$  of the matching closed (open) parenthesis.

The algorithm starts in Case 1 with  $i = 1$  ( $i \in \{1, \dots, 2n\}$ ). We associate 0 with an open parenthesis and 1 with a closed parenthesis.

- Case 1 Iterate over  $P$  until a leaf is found at position  $i$ , i.e., find an  $i$  with  $P[i] = 0 \wedge P[i + 1] = 1$ . Since we found a leaf we write a 1 as a self-delimiting number in  $A[i \dots i + 1]$ . Set  $i := i + 2$  and check if  $P[i] = 1$ . If so, move to Case 2, otherwise repeat Case 1.
- Case 2 At position  $i$  is a closing parenthesis, i.e.,  $P[i] = 1$ . In this cases we reached the end of a subtree with  $j = \text{findopen}(i)$  being the position of a corresponding open parenthesis. That means we have already computed all numbers for the whole subtree. Using an integer variable  $x$ , sum up all the self-delimiting numbers in  $A[j + 1 \dots i - 1]$ . Check if the sum  $x + 1$  exceeds  $\ell$ . If it does write 0 as a self-delimiting number in  $A[j \dots i]$  and set  $R'[j] = 1$ , otherwise write  $x + 1$  in  $A[j \dots i]$ . Note that we store only one self-delimiting number

between an open parenthesis and its matching closed parenthesis, and this number does not necessary occupy the whole space available. Hence, using `findclose` operation we jump to the end of the space that is reserved for a number and start reading the second.

After writing the number we set  $i := i + 1$ . We end the algorithm if  $i$  is out of  $P$ , otherwise we check in which case we fall next and proceed with it.

**Efficiency:** To compute the set  $R$ , we run two space-efficient DFS with  $O(n)$  bits in  $O(m + n \log^* n) = O(nk \log^* n)$  time (Lemma 1). The required  $X$ -separator and  $R$ -separator are computed in  $O(n(k \log k + \log^* n)k^4 \log^3 k)$  time (Corollary 2). We so get a balanced  $X$ -separator in  $(3^{|X|}((nk \log^* n) + (n(k \log k + \log^* n)k^4 \log^3 k))) = O(c^k n \log^* n)$  time for some constant  $c$ .

The structures  $P$ ,  $A$ , and the space-efficient DFS use  $O(n)$  bits. The size of the sets  $X_1, X_2, X_S$  and  $R$  are bound by  $O(k)$  and so use  $O(k \log n)$  bits. Corollary 2 uses  $O(n + k^2(\log k) \log n)$  bits, which is the bottleneck of our space bound.  $\square$

### 5 Tree-Decomposition Iterator Using $O(kn)$ Bits

We now introduce our iterator by showing a data structure, which we call *tree-decomposition iterator*. We think of it as an agent moving through a tree decomposition  $(T, B)$ , one node at a time in a specific order. We implement such an agent to traverse  $T$  in the order of an Euler-traversal and, when visiting some node  $w$  in  $T$ , to return the tuple  $(B(w), d_w)$  with  $d_w$  being the depth of the node  $w$ .

The tree-decomposition iterator provides the following operations:

- `init( $G, k$ )`: Initializes the structure for an undirected  $n$ -vertex graph  $G$  with treewidth  $k$ .
- `next`: Moves the agent to the next node according to an Euler-traversal and returns TRUE unless the traversal of  $T$  has already finished. In that case, it returns FALSE.
- `show`: Returns the tuple  $(B(w), d_w)$  of the node  $w$  where the agent is currently positioned.

We refer to initializing such an iterator and using it to iterate (call `show()` after every call of `next()`) over the entire tree decomposition  $(T, B)$  of a graph  $G$  as *iterating over a tree-decomposition  $(T, B)$  of  $G$* . Our goal in this section is to show that we can iterate over the bags of a tree decomposition by using  $O(kn)$  bits and  $c^k n \log n \log^* n$  time for some constant  $c$ . Recall that Reed’s algorithm computes a tree decomposition using separators. Assume we are given a separator  $S$  in  $G$ . The separator divides  $G$  into several connected components, which we have to find for a recursive call. We refer to a data structure that implements the functionality of the lemma below as a *connected-component finder*. In the next lemma we use a choice dictionary [23, 25, 30] that manages a subset  $U'$  of a universe  $U = \{1, \dots, n\}$  and provides—besides constant-time operations contains, add and remove—a linear-time iteration over  $U'$ .

**Lemma 15** *Given an undirected  $n$ -vertex  $m$ -edge graph  $G = (V, E)$  and a vertex set  $S \subseteq V$ , there is an algorithm that iterates over all connected components of  $G[V \setminus S]$ . The iterator is realized by the following operations.*

- $\text{init}(G, S)$ : Initializes the iterator.
- $\text{next}$ : Returns true exactly if there is another connected component left to iterate over. Returns false otherwise.
- $\text{show}$ : Returns the triple  $(C, p, n')$  where  $C$  is choice dictionary containing all  $n'$  vertices of a connected component, and  $p \in C$ .

The total runtime of all calls of  $\text{next}$  is  $O(n + m)$  unless  $\text{next}$  returns false, and the running time of  $\text{init}(G, S)$  as well as of  $\text{show}$  is  $O(1)$ . All operations use  $O(n)$  bits.

**Proof** The iterator is initialized by creating a bit array  $D$  to mark all vertices that are part of connected components over which we already have iterated. To hold the current state of the iterator and answer a call of  $\text{show}$  we store the triple  $(C = \emptyset, p = 0, n' = 0)$  (as defined in the lemma). Technically, to avoid modifications of the internal state of the iterator, we maintain a copy of  $C$  that we return if  $\text{show}$  is called.

If  $\text{next}$  is called, we iterate over  $V$  starting at vertex  $p$  until we either find a vertex  $p' \in V \setminus (S \cup D)$  or reach  $|V| + 1$  (out of vertices). If  $p'$  exists, we have found a connected component whose vertices are not part of  $D$ . We prepare a possible call of  $\text{show}$  by computing new values for  $(C, p, n')$  as follows: using a space-efficient BFS explore the connected component in  $G[V \setminus S]$  starting at  $p'$ . Collect all visited vertices in a choice dictionary  $C'$  as well as add them to  $D$ . Set  $C := C', p := p'$  and  $n' = |C'|$ . Finally output true. Otherwise,  $p'$  does not exist and we reached  $|V| + 1$ , set  $p' = |V| + 1$  and return false. For a call of  $\text{show}$  return  $(C, p, n')$ .

A choice dictionary can be initialized in constant time, thus  $\text{init}$  as well as  $\text{show}$  run in constant time. The operation  $\text{show}$  has to scan  $V$  for the vertex  $p'$  in  $O(n)$  time and runs a BFS in  $O(n + m)$  time to explore the connected component containing  $p'$ . However, the total time of all  $\text{next}$  operation is  $O(n + m)$  since the operation continues the scan in  $V$  from  $p$  (the last found  $p'$ ) and avoids the exploration all vertices in  $D$ , i.e., all already explored connected components. The bit array  $D$  as well as the choice dictionary  $C$  use  $O(n)$  bits. □

To implement our tree-decomposition iterator we turn Reed's recursive algorithm [37] into an iterative version. For this we use a stack structure called *record-stack* that manages a set of data structures to determine the current state of the algorithm. Informally, the record-stack allows us to pause Reed's algorithm at specific time-points and continue from the last paused point. With each recursive call of Reed's algorithm we need the following information:

- an undirected  $n_i$ -vertex graph  $G_i = (V_i, E_i)$  ( $i = 0, 1, 2, \dots$ ) of treewidth  $k$ ,
- a vertex set  $X_i$ , a separator  $S_i$ , and
- an instance  $F_i$  of the connected-component-finder data structure that iterates over the connected components of  $G[V_i \setminus S_i]$  and that outputs the vertices of each component in a bit array.

We call the combination of these elements a *record*. Although we use a record-stack structure, often we think of it to be a combination of specialized stack structures: a *subgraph-stack*, which manages to store the recursive graphs used as a parameter

for the call of Reed’s algorithm, a stack for iterating over the connected components of  $G[V \setminus S]$ , called *component-finder stack*, a stack containing the separators as bit arrays, called *S-stack*, a stack containing the vertex sets  $X$  as bit arrays, called *X-stack*. The bit arrays  $S_i$ ,  $X_i$  and  $F_i$  contain information referring to  $G_i$  and are thus of size  $O(n_i)$ . On top of  $S_i$  and  $X_i$  we create rank-select data structures. Pushing a record  $r_{\ell+1} = (G_{\ell+1}, S_{\ell+1}, X_{\ell+1}, F_{\ell+1})$  to the record-stack is equivalent to pushing each element in  $r_{\ell+1}$  to the corresponding stack (and analogous for popping).

**Lemma 16** *When a record-stack  $R$  is initialized for an undirected  $n$ -vertex graph  $G$  with treewidth  $k$  such that each subgraph  $G_i$  of  $G_0 = G$  on the subgraph-stack of  $R$  contains  $2/3$  of the vertices of  $G_{i-1}$  for  $0 < i < \ell$  and  $\ell = O(\log n)$ , then the record-stack occupies  $O(n + k \log^2 n)$  bits plus  $O(kn)$  bits for the subgraph stack.*

**Proof** Let  $m$  be the number of edges in  $G$ . We know that the size of the subgraph-stack structure is  $O(n + m)$  bits when the number of vertices of the subgraphs shrink with every push by a factor  $0 < c < 1$ . Since each subgraph of  $G_0$  has also a treewidth  $k$ , the number of edges of each subgraph is bound by  $k$  times the number of vertices. Thus, the subgraph stack uses  $O(n + m) = O(kn)$  bits.

The size of the bit arrays  $X_i$ ,  $S_i$  (including the respective rank-select structures) and the component-finder  $F_i$  is  $O(n_i)$  for  $0 \leq i \leq \ell$ . This means the total size of the stacks containing these elements is  $O(n)$  bits since they shrink in the same way as the vertex sets of the subgraphs. Storing the bag that is currently being output uses  $O(k \log n)$  bits. Thus, the size of the record-stack without the subgraph stack is  $O(n + k \log^2 n)$  bits. □

We call a tree decomposition  $(T, B)$  *balanced* if  $T$  has logarithmic height, and *binary* if  $T$  is binary. Using our space-efficient separator computation for finding a balanced  $X$ -separator we are now able to show the following theorem.

**Theorem 4** *Given an undirected  $n$ -vertex graph  $G$  with treewidth  $k$ , there exists an iterator that outputs a balanced and binary tree decomposition  $(T, B)$  of width  $8k + 6$  in Euler-traversal order using  $O(kn)$  bits and  $c^k n \log n \log^* n$  time for some constant  $c$ .*

**Proof** We implement our tree-decomposition iterator by showing `init`, `next` and `show`. We initialize the iterator for a graph  $G$  with  $n > 8k + 6$  vertices, by initializing a flag  $f = 0$ , which indicates that the agent is not yet finished, and initialize a record-stack. The record stack is initialized by first initializing its subgraph stack with a reference to  $G$  as the first graph  $G_0$ . Next, we push the empty vertex set  $X_0$  on the  $X$ -stack in form of an initial-zero bit array  $X_0$  of length  $n$ . Now, using the techniques described in Lemma 14, we find a balanced  $X_0$ -separator  $S_0$  of  $G_0$  and push it on the  $S$ -stack. Then we create a new connected-component-finder instance  $F_0$  (Lemma 15) and push  $F_0$  on the component-finder stack.

We now view our implementation of `next()`, which has the task to calculate the next bag on the fly. Since the tree  $T$  of our tree decomposition does not exist as a real structure, we only *virtually* move the agent to the next node by advancing the state of Reed’s algorithm. If  $f = 1$ , we return false (the agent can not be moved) and do not

change the state of the record-stack. Otherwise, we first virtually move the agent and then return true.

**Move to parent** If  $n_\ell \leq 8k + 6$  (we leave a leaf), we pop the record stack.

Otherwise, if the connected-component-finder instance  $F_\ell$  has iterated over all connected components and the record-stack contains more than one record, we pop it. If afterwards the record-stack contains only one record, we set  $f = 1$  (the agent moved to the root from the rightmost child; the traversal is finished).

**Move to next child** If  $F_\ell$  has not iterated over all connected components (the agent is moving to a previously untraversed node), we use  $F_\ell$  to get the next connected-component  $C$  in  $G_\ell[V_\ell \setminus S_\ell]$ , we push the vertex-induced subgraph  $G[C \cup S_\ell]$  on the subgraph stack as  $G_{\ell+1} = (V_{\ell+1}, E_{\ell+1})$  and proceed with one of the next two cases.

- If  $n_{\ell+1} \leq 8k + 6$ , we are calculating the bag of a leaf of  $T$  by setting  $B(w) = V_{\ell+1}$ . We do this by pushing a bit array with all bits set to 1 on the  $\mathcal{S}$ -stack and  $\mathcal{X}$ -stack and an empty component-finder on the component-finder stack.
- If  $n_{\ell+1} > 8k + 6$ , the agent is moving to a new internal node whose bag we are calculating as follows: we push a new bit array  $X_{\ell+1} = (X_\ell \cap V_{\ell+1}) \cup S_\ell$  on the  $\mathcal{X}$ -stack. We then find the balanced  $X_{\ell+1}$ -separator  $S_{\ell+1}$  of  $G_{\ell+1}$  and push it on the  $\mathcal{S}$ -stack. Then, create a new connected component-finder  $F_{\ell+1}$  for  $G_{\ell+1}$  and  $S_{\ell+1}$  and push it on the component-finder stack and return true.

Anytime we pop or push a new record, we call the `toptune` function of the subgraph stack to speed up the graph-access operations. This finishes our computation of `next()`.

To implement `show()`, we return the tuple  $(B(w), d_w)$  with  $B(w)$  being the current bag, and  $d_w$  being the number of records of the record-stack. The current bag is defined as  $S \cup X$ . Thus, we iterate over all elements of  $S \cup X$  via their rank-select data structures. Note that, since the subgraph  $G_\ell$  on top of the record stack is toptuned, we can return the bag as vertices of  $G_0$  or  $G_\ell$  in  $O(k)$  time.

**Efficiency** The iterator uses a record-stack structure, which occupies  $O(kn)$  bits. Since the running time of Lemma 14 is  $O(c^k n \log^* n)$ , for some constant  $c$ , and the input graphs are split in each recursion level into vertex disjoint subgraphs, the running time in each recursion level is  $O(c^k n \log^* n)$ . Summed over all  $O(\log n)$  recursion levels we get a running time of  $O(c^k n \log n \log^* n)$ .

**Make  $T$  binary** The balanced  $X$ -separator  $S$  partitions  $V \setminus S$  into any number of vertex disjoint sets between 2 and  $n$  such that no set contains more than  $2/3$  of the vertices of  $V$  (and  $X$ ). The idea is to combine these vertex sets into exactly two sets such that neither contains more than  $2/3|V|$  vertices. For this we change our usage of the connected component finder slightly. After we initialize  $F_\ell$  we also initialize two bit arrays  $C_1$  and  $C_2$  of size  $n_\ell$  each with all bits set to 0. We also store the number of bits set to 1 for each of the bit arrays as  $s_1$  and  $s_2$ , i.e., the number of vertices contained in them (initially 0). We now want to collect the vertices of all connected components of  $G_\ell[V_\ell \setminus S_\ell]$  in  $C_1$  and  $C_2$ . While there are still connected components to be returned by  $F_\ell$ , this is done by obtaining the size of the next connected component via  $F_\ell$  as  $s$ . If  $s_1 + s \leq 2/3|V_\ell|$ , we collect the next connected component in  $C_1$  and increment

$s_1 = s_1 + s$ . Otherwise, we do the same but for  $C_2$  and  $s_2$ . Doing this until all connected components are found results in  $C_1$  and  $C_2$  to contain all connected components of  $G_\ell[V_\ell \setminus S_\ell]$ . For  $(C_1, C_2)$  we implement a function that returns  $C_1$  if it was not yet returned, or  $C_2$  if it was not yet returned, or null otherwise. We store  $(C_1, C_2)$  with the respective functions on the connected component finder stack (instead of  $F_\ell$ ). Any time we do this during our iterator, the graph  $G_\ell$  is toptuned, resulting in constant time graph access operations. The previous runtime and space bounds still hold.  $\square$

Often it is needed to access the subgraph  $G[B(w)]$  induced by a bag  $B(w)$  of a tree decomposition  $(T, B)$  for further computations. We call such a subgraph *bag-induced*. For this we show the following:

**Lemma 17** *Given an undirected  $n$ -vertex graph  $G$  with treewidth  $k$  and an iterator  $A : \mathcal{G} \rightarrow (T, B)$  that iterates over a balanced tree decomposition of width  $O(k)$  we can additionally output the bag-induced subgraphs using  $O(k^2 \log n)$  bits additional space and  $O(kn \log n)$  additional time.*

**Proof** To obtain the edges we use a bit matrix  $M_\ell$  of size  $O(k^2)$  whose bit at index  $[v][u]$  is set to 1 exactly if there exists an edge  $\{v, u\}$  in  $G[B(w)_\ell]$ . To quickly find the edges in  $G[B(w)_\ell]$  we use  $M_\ell$  together with rank-select data structures on  $S_\ell$  and  $X_\ell$  that allow us to map the vertices in  $B(w)_\ell = S_\ell \cup X_\ell$ , names as vertices of  $G_\ell$ , to  $\{1, \dots, k\}$ .

We create  $M_\ell$  anytime a new record  $r_\ell$  is pushed on the record stack, and anytime  $r_\ell$  is popped, we throw away  $M_\ell$ . For reasons of performance relevant in Sect. 7, we want to avoid accessing vertices multiple times in a graph. Hence, when we push a bit matrix  $M_\ell$ , we first use  $M_{\ell-1}$  to initialize edges that were contained in  $B(w)_{\ell-1}$  and are still contained in  $B(w)_\ell$ . To obtain the edges of a vertex  $v$  in  $B(w)_\ell$ , we can iterate over all the edges of  $v$  in  $G_\ell$  and check if the opposite endpoint is in  $B(w)_\ell$ . However, by using definition (TD2) of a tree decomposition, it suffice to iterate only over the edges of such a vertex once, i.e., the first time they are contained in a bag to create  $M_\ell$ .

**Efficiency** We can see that we store at most  $O(\log n)$  matrices this way since the record stack contains  $O(\log n)$  records (i.e., the height of the treedecomposition). Storing all bit matrices uses  $O(k^2 \log n)$  bits and initializing all bit matrices takes  $O(kn \log n)$  time, including initializing and storing the rank-select structures if not already present.  $\square$

We conclude the section with a remark on the output scheme of our iterator. The specific order of an Euler-traversal encompasses many other orders of tree traversal such as pre-order, in-order or post-order. To achieve these orders we simply filter the output of our iterator, i.e., skip some output values.

## 6 Modifying the Record Stack to Work with $O(n + k^2 \log n)$ Bits

The space requirements of the record stack used by the iterator shown in Theorem 4 is  $O(kn)$  bits. Our goal in Sect. 7 is to reduce it to  $O(n)$  bits. The bottleneck here

is the record stack. Assume that, for  $\ell = O(\log n)$ , a graph  $G_\ell = (V_\ell, E_\ell)$  with  $n_\ell$  vertices is on top of the record stack. When considering the record  $r_\ell$  on top of the record stack we see that most structures use  $O(n_\ell)$  bits: a separator  $S_\ell$ , a vertex set  $X_\ell$  and a connected component finder  $F_\ell$ . The only structure that uses more space is the subgraph stack, which uses  $O(kn_\ell)$  bits. This is due to the storage of the edge set  $E_\ell$  using  $O(kn_\ell)$  bits. The strategy we want to pursue is to store only the vertices of the subgraphs (but not the edges) such that the space requirement of the subgraph stack is  $O(n)$  bits. We call such a subgraph stack a *minimal subgraph stack*. In the following we always assume that the number of subgraphs on the minimal subgraph stack is  $O(\log n)$  and that the subgraphs shrink by a constant factor. This is in particular the case for the subgraphs generated by Reed's algorithm.

In the following we make a distinction between *complete* and *incomplete* vertices. Only complete vertices have all their original edges, i.e., they have the same degree in the original graph as they do in the subgraph. Note that the number of incomplete vertices in each subgraph is  $O(k)$ , which follows directly from the separator size. To clarify, a vertex in the subgraph  $G_\ell$  on top of the subgraph stack is incomplete exactly if it is contained in a separator of the parent graph  $G_{\ell-1}$ .

**Lemma 18** *Assume that we are given an undirected  $n$ -vertex graph  $G = (V, E)$  with treewidth  $k$  as well as a toptuned minimal subgraph stack  $(G_0 = G, \dots, G_\ell)$  with  $\ell \in O(\log n)$ . Assume further that each graph  $G_i$  ( $1 \leq i \leq \ell$ ) has  $n_i$  vertices,  $m_i$  edges and contains  $O(k)$  incomplete vertices. The modified subgraph stack can be realized with  $O(n + k^2 \log n)$  bits and allows us to push an  $n_{\ell+1}$ -vertex graph  $G_{\ell+1}$  on top of a minimal subgraph stack in  $O(k^2 n_\ell \log^* \ell)$  time. The resulting graph interface allows us to access the adjacency array of the complete vertices in constant time whereas an iteration over the adjacency list of an incomplete vertex runs in  $O(m_\ell)$  time.*

**Proof** Recall that the subgraph stack considers every edge as a pair of directed arcs. Let  $\phi$  be the vertex translation between  $G_i$  and  $G_0$ . Each complete vertex of  $G_i = (V_i, E_i)$  has the same degree in both  $G_i$  and  $G_0$ . Thus, to iterate over all arcs of a complete vertex  $v \in V_i$ , iterate over every arc  $(\phi(v), u)$  of  $\phi(v)$  and return the arc  $(v, \phi^{-1}(u))$ . For a complete vertex  $v$  we can use the adjacency array of  $v$ . To iterate over the arcs of an incomplete vertex  $v$  (via support of adjacency lists of  $v$ ), we differ two cases: (1) the arcs to a complete vertex and (2) the arcs to another incomplete vertex. To iterate over all arcs of (1) we iterate over all complete vertices  $u$  of  $G_i$  and check in  $G_0$  if  $\phi(u)$  has an edge to  $\phi(v)$ . If it does,  $\{v, u\}$  is an edge of  $v$ . Thus, the iteration over all arcs incident to an incomplete vertex of (1) runs in  $O(m_i)$  time.

For the arcs according to (2), we use matrices  $M_i$  storing the edges in between incomplete vertices. To build the matrices we proceed as follows. Whenever a new graph  $G_{\ell+1}$  is pushed on the subgraph stack, we create a bit matrix  $M_{\ell+1}$  of size  $k^2$  and a rank-select data structure  $I_{\ell+1}$  of size  $n_{\ell+1}$  with  $I_{\ell+1}[v] = 1$  exactly if  $v$  is incomplete.  $M_{\ell+1}$  is used to store the information if  $G_{\ell+1}$  contains an edge  $\{u', v'\}$  between any two incomplete vertices  $u'$  and  $v'$  of  $G_{\ell+1}$ , which is the case exactly if  $M_{\ell+1}[I_{\ell+1}.\text{rank}(u')][I_{\ell+1}.\text{rank}(v')] = 1$  and  $M_{\ell+1}[I_{\ell+1}.\text{rank}(v')][I_{\ell+1}.\text{rank}(u')] = 1$ .

First, we initialize  $M_{\ell+1}$  to contain only 0 for all bits. Then we use  $M_\ell$  to find edges between incomplete vertices of  $G_\ell$  and set the respective bits in  $M_{\ell+1}$  to 1 if

those incomplete vertices are still contained in  $G_{\ell+1}$  (if  $\ell = 0$ , we set all bits to 0). Afterwards we are able to find edges between incomplete vertices that are already incomplete in the previous graph. We need to update  $M_{\ell+1}$  to contain the information of the edges between the vertices that are complete in  $G_\ell$ , but are not complete in  $G_{\ell+1}$ . Since they are complete in  $G_\ell$ , we can simply iterate over all complete edges  $e$  of  $G_\ell$  in  $O(km_\ell)$  time and check if both endpoints of  $e$  are incomplete in  $G_{\ell+1}$  via  $I_{\ell+1}$ . If so, we set the respective bits in  $M_{\ell+1}$  to 1.

**Efficiency** Queries on  $M_i$  ( $1 \leq i \leq \ell$ ) allow us to iterate over all arcs of (2) of an incomplete vertex in  $O(k)$  time. This results in a combined runtime of  $O(m_\ell + k) = O(m_\ell)$  by ignoring zero-degree vertices in  $M_\ell$ .

Storing all bit matrices  $M_i$  uses  $O(k^2\ell) = O(k^2 \log n)$  bits and the space used by the rank-select structures is negligible.

The adjacency lists are realized by storing a pointer for each vertex. This uses negligible additional  $\Theta(k \log n)$  bits for implementing the interface. Our modified subgraph stack uses  $O(n + k^2 \log n)$  bits. □

The last lemma allows us to store all recursive instances of Reed’s algorithm with  $O(n + k^2 \log n)$  bits. We use the result in the next section to show our first  $O(n)$ -bit iterator to output a tree decomposition on graphs of small treewidth.

### 7 Tree-Decomposition Iterator Using $O(n)$ Bits for $k = O(n^{1/2-\epsilon})$

By combining the  $O(kn)$ -bit iterator of Theorem 4 with the modified record stack of Lemma 18 we can further reduce the space to  $O(n)$  bits for a sufficiently small  $k$ . Recall that the only structure using more than  $O(n)$  bits in the proof of Theorem 4 was the subgraph stack whose space is stated in Lemma 16. This allows us to show the following theorem.

**Theorem 5** *Given an undirected  $n$ -vertex graph  $G$  with treewidth  $k$ , there is an iterator that outputs a balanced binary tree decomposition  $(T, B)$  of width  $8k + 6$  in Euler-traversal order using  $O(n + k^2 \log^2 n)$  bits and  $c^k n \log n \log^* n$  time for some constant  $c$ . For  $k = O(n^{1/2-\epsilon})$  with an arbitrary  $\epsilon > 0$ , our space consumption is  $O(n)$  bits.*

**Proof** Recall that the tree decomposition iterator of Theorem 4 uses the algorithm of Theorem 3 to find  $k$  vertex-disjoint paths for the construction of the separators.

With the construction of the  $k$  vertex-disjoint paths we have to compute and access path storage schemes. To avoid querying the neighborhood of an incomplete vertex  $v$  several times by querying the region with  $v$  several times, we add the  $O(k)$  incomplete vertices to the boundary vertices when we build a path data structure. This neither increases our asymptotical time nor our total space bounds. Moreover, by adding the  $O(k)$  incomplete vertices  $v$  to the boundary vertices and storing  $\text{prev}(v)/\text{next}(v)$ , we avoid running a BFS on incomplete vertices and searching for the predecessor and successor of  $v$ . If we afterwards construct a path, we iterate over the adjacency lists of incomplete vertices only once—even if regions are constructed several times.



The construction of a vertex-disjoint path involves  $O(\log k)$  reroutings. Recall that each rerouting extends a clean area by traversing over a path  $P^*$  (two DFS) and then searches “backwards” within the area with two further DFS. Altogether, a rerouting can be done with  $O(1)$  DFS runs and  $O(\log k)$  DFS runs over all reroutings for one path  $P^*$ .

To construct a balanced  $X$ -separator a set  $R$  is constructed by another DFS run. However, this iterates over the neighbors of each vertex only once. To sum up a balanced  $X$ -separator can be constructed by  $O(\tilde{c}^k)$  DFS runs for some constant  $\tilde{c}$ .

Given an  $n_i$ -vertex,  $m_i$ -edge graph  $G_i$  stored in a minimal subgraph stack, we can run the space-efficient DFS of Lemma 1 in  $O(km_i + n_i \log^* n_i) = O(k^2 n_i \log^* n_i)$  time and  $O(n + k \log n)$  bits, which is a factor of  $O(k)$  slower than the DFS of Theorem 2. Thus, a balanced  $X$ -separator can be constructed in  $O(\tilde{c}^k n_i \log^* n_i)$  time for some constant  $\tilde{c}$ , which is the same time as stated in the proof of Theorem 4 (the constant  $\tilde{c}$  here is larger). Concerning the space consumption note that we have a record-stack of size  $O(n + k \log^2 n)$  bits by Lemma 16 by using our minimal subgraph stack of  $O(n + k^2 \log n)$  bits. The  $O(k)$  extra values for `prev` and `next` are negligible. Finally note that we can search a separator with  $O(n + k^2(\log k) \log n)$  bits by Corollary 2.  $\square$

We next combine the theorem above with a recent tree-decomposition algorithm by Bodlaender et al. [11] as follows. The algorithm by Bodlaender et al. finds a tree decomposition for a given  $n$ -vertex graph  $G$  of treewidth  $k$  in  $b^k n$  time for some constant  $b$  [11]. The resulting tree decomposition has a width of  $5k + 4$ . The general strategy pursued by them is to first compute a tree decomposition of large width and then use dynamic programming on that tree decomposition to obtain the final tree decomposition of width  $5k + 4$ . For an overview of the construction, we refer to [11, p. 3]. The final tree decomposition is balanced due to the fact that its construction uses balanced  $X$ -separators at every second level, alternating between an 8/9-balanced and an unbalanced  $X$ -separator [11, p. 26]. Further details of the construction of different kinds of the final tree decomposition can be found in [11, p. 20, and p. 39]. Since its runtime is bounded by  $b^k n$ , it can write at most  $b^k n$  words and thus has a space requirement of at most  $b^k n \log n$  bits.

Our following idea is to use a hybrid approach to improve the runtime of our iterator. We first run our iterator (Theorem 5). Once the height of the record-stack of our tree-decomposition iterator is equal to  $z = b^k \log \log n$ , the call of `next()` uses an unbalanced  $X$ -separator  $S^*$ . This ensures that the size of the bag is at most  $4k + 2$  instead of  $8k + 6$ . (We later add all vertices in the bag to all following bags.) Note that using a single unbalanced  $X$ -separator  $S^*$  on all root-to-leaf paths of our computed tree decomposition increases the height of the tree decomposition only by one. A following call of `next()` toptunes the graph  $G_\ell$  and then uses Bodlaender et al.’s linear-time tree-decomposition algorithm [11] to calculate a tree decomposition  $(T', B')$  of an  $n_\ell$  vertex subgraph  $G_\ell$ , which we then turn by folklore techniques into a binary tree decomposition  $(T'', B'')$  by neither increasing the asymptotic size nor the width of the tree decomposition. In detail, this is done by repeatedly replacing all nodes  $w$  with more than two children by a node  $w_0$  with two children  $w_1$  and  $w_2$ , with  $B(w_0) = B(w_1) = B(w_2) = B(w)$ , followed by adding the original children

of  $w$  to  $w_1$  and  $w_2$ , alternating between them both. To ensure property (TD2) of a tree decomposition, we add the vertices in  $S^*$  to all bags of  $(T'', B'')$ . We so get a tree decomposition of the width  $(5k + 4) + (4k + 2) = 9k + 6$ .

Since  $G_\ell$  contains  $n_\ell = O(n/2^z) = O(n/(b^k \log n))$  vertices, the space usage of the linear-time tree-decomposition algorithm is  $b^k n_\ell \log n = O(n)$  bits. The runtime of the algorithm is bounded by  $b^k n_\ell$ . Once we obtain  $(T', B')$ , we also need to transform each bag  $b'$  of  $B'$  since  $B'$  contains mappings in relation to  $G'$ , but we want them to contain mappings in relation to  $G$ . This can be done in negligible time since  $G'$  was toptuned before. We then initialize a tree-decomposition iterator  $I$  for  $(T', B')$  as described in the beginning of Sect. 5. Now, as long as  $I'$  has not finished its traversal of  $(T', B')$ , a call to `next` on  $I$  is equal to a call to `next` on  $I'$ . Similarly, a call to `show` on  $I$  now returns the tuple  $(B'(w), d_w)$  with  $d_w$  being the depth of  $w$  in  $T'$  plus the size of the record stack of  $I$ . Once iterator  $I'$  is finished, we throw away  $(T', B')$ . Then, the operations of `next` and `show` work normally on  $I$  until the size of the record stack again is  $O(b^k \log \log n)$  or until the iteration is finished. Since we use our iterator only to recursion depth  $z$ , our algorithm runs in  $a^k n (\log^* n)^z$  time for some constant  $a$ . The total runtime is  $a^k n (\log^* n) (b^k \log \log n) + b^k n = c^k n \log \log n \log^* n$  for some constant  $c$ .

**Corollary 3** *There is an iterator to output a balanced binary tree decomposition  $(T, B)$  of width  $9k + 6$  for an  $n$ -vertex  $G = (V, E)$  with treewidth  $k$  in Euler-traversal order in  $c^k n \log \log n \log^* n$  time for some constant  $c$  using  $O(n + k^2 \log^2 n)$  bits. For  $k = O(n^{1/2-\epsilon})$  and an arbitrary  $\epsilon > 0$ , the space consumption is  $O(n)$  bits.*

If we try to run our iterator from the last corollary on a graph that has a treewidth greater than  $k$ , then either the computation of a vertex separator or the computation of Bodleander et al.’s algorithm for finding a tree decomposition [11] fails. In both cases, our iterator stops and we can output that the treewidth of  $G$  is larger than  $k$ .

## 8 Applications

Courcelle [16] showed that all monadic-second-order (MSO) problems can be solved in polynomial time on graphs with treewidth bounded by a constant, and Elberfeld et al. [19] showed that the same is possible when using logarithmic space. As briefly mentioned in the introduction, due to Banerjee et al. [4] there exists a framework that takes as input a graph  $G$  and a tree decomposition  $(T, B)$  and computes a solution set for a given MSO problem. Their framework allows a tradeoff between the runtime and space usage. In particular, they showed the following theorem.

**Theorem 6** [4] *Let  $G$  be an  $n$ -vertex graph given with a tree decomposition  $(T = (V_T, E_T), B)$  of width  $k$ . Then a solution set for all weighted MSO problems can be found in  $O(\tau(k) \cdot n^{2+(2/\log p)})$  time and  $O(\tau(k) \cdot p \log_p n)$  variables, i.e.,  $O(\tau(k) \cdot p(\log_p n) \log n)$  bits, for any parameter  $2 \leq p \leq n$  and a problem dependent function  $\tau$  depending on the MSO formula.*

As we are able to provide the tree decomposition interface required by their framework in  $O(n)$  bits (for graphs of small enough treewidth) via Corollary 3, this directly implies the following theorem.

**Theorem 7** *Let  $G$  be an  $n$ -vertex graph with treewidth  $k$ . Then a solution set for all weighted MSO problems can be found in time  $O(\tau(k) \cdot n^{2+(2/\log p)})$  time and  $O(n + \tau(k) \cdot p(\log_p n) \log n)$  bits for any parameter  $2 \leq p \leq n$  and a problem dependent function  $\tau$  depending on the MSO formula.*

In the following we outline a faster  $O(n)$ -bit approach for vertex cover that can easily be generalized to other problems. We know from [17, Theorem 7.9] that there is an algorithm that solves all problems mentioned in Theorem 8 on an  $n$ -vertex graph with treewidth  $k$  in  $c^k n$  time for some constant  $c$  when a tree decomposition with approximation-ratio  $O(1)$  is given. The general strategy used for solving these problems is almost identical. First, traverse the tree decomposition bottom-up and compute a table for each node  $w$ . The table stores the size of all best possible solutions in the graph induced by all bags belonging to nodes below  $w$  under certain conditions for the vertices in bag  $B(w)$ . E.g., for VERTEX COVER the table contains  $2^{k+1}$  solutions ( $v \in B(w)$  does belong or does not belong to the solution) and for DOMINATING SET it contains  $3^{k+1}$  solutions (one additionally differs, if a vertex is already dominated or not). We only consider problems whose table has at most  $c^k$  solutions for some constant  $c$ . For each possible solution, the table stores the size of the solution and thus uses  $O(c^k \log n)$  bits. After the bottom-up traversal, the minimal/maximal solution size in the table at the root is the solution for the minimization/maximization problem, respectively. An optimal solution set can be obtained in a top-down traversal by using the tables.

It is clear that, for large  $k$ , we can not store all tables when trying to use  $O(n)$  bits. Our strategy is to store the tables only for the nodes on a single root-leaf path of the tree decomposition and for nodes with a depth less than some threshold value. The other tables are recomputed during the construction of the solution set. For a balanced tree decomposition this results in  $O(c^k \log^2 n)$  bits storing  $O(\log n)$  tables for vertices on a root-leaf path each of size  $O(c^k \log n)$  bits. Using this strategy we have all information to use the standard bottom-up traversal to compute the size of the solution for the given problem for  $G$ . To obtain an optimal solution set we need a balanced and binary tree decomposition with a constant-factor approximation of the treewidth.

We conclude this section by giving a list of problems that can be solved with the same asymptotic time and space bound.

A vertex cover of a graph  $G = (V, E)$  is a set of vertices  $C \subseteq V$  such that, for each edge  $\{u, v\} \in E$ ,  $u \in C \vee v \in C$  holds. For graphs with a small treewidth, one can find a minimum vertex cover by first computing a tree decomposition of the graph and then, using dynamic programming, calculate a minimal vertex cover. We start to sketch the standard approach.

Let  $(T, B)$  be a tree decomposition of width  $O(k)$  of an undirected graph  $G$  with treewidth  $k$ . Now, iterate over  $T$  in Euler-traversal order and, if a node  $w$  is visited for the first time, calculate and store in a table  $\mathcal{T}_w$  all possible solutions of the vertex

cover problem for  $G[B(w)]$ . Also store the value of each solution, which is equal to the number of vertices used for the cover. If the solution is not valid, store  $\infty$  instead.

When visiting a node  $w$  and  $\mathcal{T}_w$  already exists, we update  $\mathcal{T}_w$  by using  $\mathcal{T}_{w'}$  with  $w'$  being the node visited during the Euler-traversal right before  $w$  ( $w'$  is a child of  $w$ ). The update process is done by comparing each solution  $s$  in  $\mathcal{T}_w$  with each overlapping solution in  $\mathcal{T}_{w'}$ . A solution  $s' \in \mathcal{T}_{w'}$  is chosen if it has the smallest value among overlapping solution. The value of  $s'$  is added to the value of  $s$ , and the two solutions are linked with a pointer structure. Two solutions  $s$  and  $s'$  are overlapping exactly if, for each  $v \in B(w) \cap B(w')$ ,  $(v \in s \wedge v \in s') \vee (v \notin s \wedge v \notin s')$  is true. Once the Euler-traversal is finished, the table  $\mathcal{T}_r$ , with  $r$  being the root of  $T$ , contains the size of the minimum vertex cover  $C$  of  $G$  as the smallest value of all solutions. This is the first step of the algorithm.

The second step is obtaining  $C$ , which is done by traversing top-down through all tables with the help of the pointer structures, starting at the solution with the smallest value in  $\mathcal{T}_r$ , and adding the vertices used by the solutions to the initially empty set  $C$  if they are not yet contained in  $C$ . For simplicity, we first only focus on VERTEX COVER, but use a problem specific constant  $\lambda$  for the size of the tables and runtime of solving subproblems. For VERTEX COVER  $\lambda = 2$ . This allows us an easy generalization step to other problems subsequently. In the following the default base of log is 2.

**Lemma 19** *Given an  $n$ -vertex graph  $G$  with treewidth  $k \leq \log_\lambda n - 2 \log_\lambda \log n$  the size of an optimal VERTEX COVER  $C$  of  $G$  can be computed with  $O(n)$  bits in  $c^k n \log \log n \log^* n$  time for some constant  $c$ .*

**Proof** For an  $n$ -vertex graph  $G$  with treewidth  $k$  and a given tree decomposition  $(T, B)$  of width  $k' = O(k)$  the runtime of the algorithm is  $O(2^{k'} n)$ . A table  $\mathcal{T}_w$  constructed for a bag  $B(w)$  consists of a bit array of size  $O(k')$  for each of the  $2^{k'+1}$  possible solutions, and their respective values and pointer structures. This uses  $O(2^{k'+1}(k + \log n)) = O(\lambda^k \log n)$  bits per table for some problem specific constant  $a$ , which is two for VERTEX COVER. Thus, storing the tables for the entire tree decomposition uses  $O(\lambda^k n \log n)$  bits. Our goal is to obtain the optimal vertex cover using only  $O(n)$  bits for both the tree decomposition  $(T, B)$  and the storage of the tables. For obtaining only the size of  $C$ , i.e., the first step of the algorithm, we only need to store the tables for the current root-node path of the tree decomposition iterator. The reason is that once a table has been used to update its parent table it is only needed for later obtaining the final cover via the pointer structures. We can iterate over a balanced binary tree decomposition of width  $O(k)$  in  $\tilde{c}^k n \log \log n \log^* n$  time using  $O(n)$  bits (Theorem 3) for some constant  $\tilde{c}$ . To obtain the bag-induced subgraphs we use Lemma 17. We have to store  $O(\log n)$  tables, which results in  $O(\lambda^k \log^2 n)$  bits used, which for  $k \leq \log_\lambda n - 2 \log_\lambda \log n$  equals  $O(n)$  bits ( $\lambda^k \log^2 n \leq n \Rightarrow \lambda^k \leq n \log^{-2} n \Rightarrow k \leq \log_\lambda n - 2 \log_\lambda \log n$ ). Initializing and updating all tables can be done in  $O(\lambda^k n)$  time. □

To obtain the final vertex cover we need access to all tables and bags they have been initially created for. We now use the previous lemma with modifications. Our idea is to fix some  $\ell \in \mathbb{N}$  and to use partial tree decompositions of depths  $\ell$  rooted at every  $\ell$ th node of a root-leaf path. For this, let us define  $\text{ptd}_{G, \ell}(w)$  to be the partial tree decomposition  $(T', B)$  where  $T'$  is a subtree of  $T$  with root  $w$  and depths  $\ell$ .

**Lemma 20** *Assume that we are given an  $n$ -vertex graph  $G$  with treewidth  $k \leq \log_\lambda n - 3 \log_\lambda \log n = c' \log n$  for some constant  $0 < c' < 1$ . Then we can calculate the optimal vertex cover  $C$  of  $G$  in  $c^k n \log n \log^* n$  time using  $O(n)$  bits for some constant  $c$ .*

**Proof** By iterating over  $(T, B)$  we first compute the tables and pointer structures for  $(T', B) = \text{ptd}_{G,\ell}(r)$  where  $r$  is the root of the tree decomposition of  $G$  and where  $\ell = \log \log n$ . Thus, the partial tree  $T'$  consists of  $O(\log n)$  nodes, each with a table of  $O(\lambda^k \log n)$  bits. (Recall that for VERTEX COVER  $\lambda = 2$ .)

We then start to follow the pointer structures starting from  $r$ . When we arrive at a node  $w$  having a table where the pointer structure is invalid (because the next table does not exist) we build  $\text{ptd}_{G,\ell}(w)$ . Afterwards, we can continue to follow the pointer structures since the next tables now exist. We repeat to build the partial tree decompositions anytime we try to follow a pointer that is invalid until we arrive at a leaf (at which point we backtrack). When we have built and processed all tables of a subtree with root  $w$ , we can throw away all tables of  $\text{ptd}_{G,\ell}(w)$ . In other words, we have to store tables for only  $O(\log n / \ell)$  partial trees. Note that each partial tree uses  $O(\log n)$  tables of  $O(\lambda^k \log^2 n)$  bits in total. Summed over all partial trees on a root-leaf path, we need to store  $O(\lambda^k \log^3 n)$  bits. For  $k \leq \log_\lambda n - 3 \log_\lambda \log n$  this uses  $O(n)$  bits ( $\lambda^k \frac{\log^3 n}{\log \log n} \leq n \Rightarrow \lambda^k \leq n \frac{\log \log n}{\log^3 n} \Rightarrow k \leq \log_\lambda n - 3 \log_\lambda \log n + \log_\lambda \log \log n$ ). It remains to show the impact on the runtime. Anytime we want to obtain the tables of  $\text{ptd}_{G,\ell}(w)$ , we need to compute the tables of the subtree rooted by  $w$ . Thus partial subtrees  $\text{ptd}_{G,\ell}(w)$  with deepest nodes  $w$  need to be calculated  $O(\log n / \log \log n)$ -times, the partial subtrees above  $(O(\log n / \log \log n) - 1)$ -times and so forth. This can be thought of as iterating over the tree decomposition  $(T, B)$  of  $G$  for  $O(\log n / \log \log n)$  times. In other words, we run the algorithm of Theorem 19,  $O(\log n / \log \log n)$  times.  $\square$

We finally present our last theorem.

**Theorem 8** *Let  $G$  be an  $n$ -vertex graph with treewidth  $k \leq c' \log n$  for some constant  $0 < c' < 1$ . Using  $O(n)$  bits and  $c^k n \log n \log^* n$  time for some constant  $c$  we can solve the following problems: VERTEX COVER, INDEPENDENT SET, DOMINATING SET, MAXCUT and  $q$ -COLORING.*

**Proof** As in the proof of Lemma 19  $k' = O(k)$  is the width of the tree decomposition. To change the previous two proofs to the different problems mentioned in Theorem 8 the only change is in the computation of the tables and the size of the tables. Thus, we have to take another value for the problem specific constant  $\lambda$  of the Lemmas 19 and 20. For VERTEX COVER, INDEPENDENT SET and MAX CUT the tables contain  $2^{k'+1}$  possible solutions and thus  $\lambda = 2$ . For DOMINATING SET they contain  $3^{k'+1}$  and  $q$ -COLORING they contain  $q^{k'+1}$  possible solutions, so  $\lambda = 3$  and  $\lambda = q$ , respectively. For further details see [10].  $\square$

**Acknowledgements** We thank the anonymous reviewers for their careful reading of our manuscript and their helpful comments.

**Funding** Open Access funding enabled and organized by Projekt DEAL. Johannes Meintrup and Andrej Sajenko: Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—379157101.

**Availability of Data and Materials** Not applicable.

## Declarations

**Conflict of interest** Not applicable.

**Code Availability** Not applicable.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. Ahuja, R.K., Magnanti, T.L., Orlin, J.B.: Network Flows: Theory, Algorithms and Applications. Prentice Hall, Hoboken (1993)
2. Amir, E.: Approximation algorithms for treewidth. *Algorithmica* **56**(4), 448–479 (2010). <https://doi.org/10.1007/s00453-008-9180-4>
3. Asano, T., Izumi, T., Kiyomi, M., Konagaya, M., Ono, H., Otachi, Y., Schweitzer, P., Tarui, J., Uehara, R.: Depth-first search using  $o(n)$  bits. In: Proceedings of 25th International Symposium on Algorithms and Computation (ISAAC 2014). LNCS, vol. 8889, pp. 553–564. Springer (2014). [https://doi.org/10.1007/978-3-319-13075-0\\_44](https://doi.org/10.1007/978-3-319-13075-0_44)
4. Banerjee, N., Chakraborty, S., Raman, V., Roy, S., Saurabh, S.: Time-space tradeoffs for dynamic programming algorithms in trees and bounded treewidth graphs. In: Proceedings of 21st International Conference on Computing and Combinatorics (COCOON 2015). LNCS, vol. 9198, pp. 349–360. Springer (2015). [https://doi.org/10.1007/978-3-319-21398-9\\_28](https://doi.org/10.1007/978-3-319-21398-9_28)
5. Banerjee, N., Chakraborty, S., Raman, V., Satti, S.R.: Space efficient linear time algorithms for BFS, DFS and applications. *Theory Comput. Syst.* **62**(8), 1736–1762 (2018). <https://doi.org/10.1007/s00224-017-9841-2>
6. Barbay, J., Aleardi, L.C., He, M., Munro, J.I.: Succinct representation of labeled graphs. *Algorithmica* **62**(1), 224–257 (2012). <https://doi.org/10.1007/s00453-010-9452-7>
7. Baumann, T., Hagerup, T.: Rank-select indices without tears. In: Proceedings of 16th International Symposium on Algorithms and Data Structures (WADS 2019). LNCS, vol. 11646, pp. 85–98. Springer (2019). [https://doi.org/10.1007/978-3-030-24766-9\\_7](https://doi.org/10.1007/978-3-030-24766-9_7)
8. Bodlaender, H., Gilbert, J., Hafsteinsson, H., Kloks, T.: Approximating treewidth, pathwidth, frontsize, and shortest elimination tree. *J. Algorithms* **18**(2), 238–255 (1995). <https://doi.org/10.1006/jagm.1995.1009>
9. Bodlaender, H.L.: A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.* **25**(6), 1305–1317 (1996). <https://doi.org/10.1137/S0097539793251219>
10. Bodlaender, H.L., Bonsma, P., Lokshtanov, D.: The fine details of fast dynamic programming over tree decompositions. In: Parameterized and Exact Computation, pp. 41–53. Springer (2013)
11. Bodlaender, H.L., Drange, P.G., Dregi, M.S., Fomin, F.V., Lokshtanov, D., Pilipczuk, M.: A  $c^k n$  5-approximation algorithm for treewidth. *SIAM J. Comput.* **45**(2), 317–378 (2016). <https://doi.org/10.1137/130947374>
12. Brodnik, A., Munro, J.I.: Membership in constant time and almost-minimum space. *SIAM J. Comput.* **28**(5), 1627–1640 (1999). <https://doi.org/10.1137/S0097539795294165>

13. Chatterjee, K., Goharshady, A.K., Goharshady, E.K.: The treewidth of smart contracts. In: Proceedings of 34th ACM/SIGAPP Symposium on Applied Computing (SAC 2019), pp. 400–408. ACM (2019). <https://doi.org/10.1145/3297280.3297322>
14. Choudhari, J., Gupta, M., Sharma, S.: Nearly optimal space efficient algorithm for depth first search. CoRR [arXiv:1810.07259](https://arxiv.org/abs/1810.07259) (2018)
15. Clark, D.R.: Compact pat trees. Ph.D. thesis, Waterloo, Ont., Canada, Canada (1998). UMI Order No. GAXNQ-21335
16. Courcelle, B.: The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. Inf. Comput. **85**, 12–75 (1990)
17. Cygan, M., Fomin, F.V., Kowalik, L., Lokshtanov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: Parameterized Algorithms. Springer, Berlin (2015). <https://doi.org/10.1007/978-3-319-21275-3>
18. Dumitrescu, E.F., Fisher, A.L., Goodrich, T.D., Humble, T.S., Sullivan, B.D., Wright, A.L.: Benchmarking treewidth as a practical component of tensor-network-based quantum simulation. CoRR [arXiv:1807.04599](https://arxiv.org/abs/1807.04599) (2018)
19. Elberfeld, M., Jakoby, A., Tantau, T.: Logspace versions of the theorems of Bodlaender and Courcelle. In: Electronic Colloquium on Computational Complexity (ECCC 2010), vol. 17, p. 62 (2010)
20. Elmasry, A., Hagerup, T., Kammer, F.: Space-efficient basic graph algorithms. In: 32nd International Symposium on Theoretical Aspects of Computer Science, (STACS 2015). LIPIcs, vol. 30, pp. 288–301. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2015). <https://doi.org/10.4230/LIPIcs.STACS.2015.288>
21. Fafianie, S., Kratsch, S.: Streaming kernelization. In: Proceedings of 39th International Symposium on Mathematical Foundations of Computer Science (MFCS 2014), pp. 275–286 (2014). [https://doi.org/10.1007/978-3-662-44465-8\\_24](https://doi.org/10.1007/978-3-662-44465-8_24)
22. Feige, U., Hajiaghayi, M.T., Lee, J.R.: Improved approximation algorithms for minimum-weight vertex separators. In: Proceedings of 37th ACM Symposium on Theory of Computing (STOC 2005), pp. 563–572. ACM (2005). <https://doi.org/10.1145/1060590.1060674>
23. Hagerup, T.: Small uncolored and colored choice dictionaries. CoRR [arXiv:1809.07661](https://arxiv.org/abs/1809.07661) (2018)
24. Hagerup, T.: Space-efficient DFS and applications to connectivity problems: simpler, leaner, faster. Algorithmica **82**(4), 1033–1056 (2020). <https://doi.org/10.1007/s00453-019-00629-x>
25. Hagerup, T., Kammer, F.: Succinct choice dictionaries. CoRR [arXiv:1604.06058](https://arxiv.org/abs/1604.06058) (2016)
26. Hagerup, T., Kammer, F.: On-the-fly array initialization in less space. In: Proceedings of 28th International Symposium on Algorithms and Computation (ISAAC 2017). LIPIcs, vol. 92, pp. 44:1–44:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2017). <https://doi.org/10.4230/LIPIcs.ISAAC.2017.44>
27. Hagerup, T., Kammer, F., Laudahn, M.: Space-efficient Euler partition and bipartite edge coloring. Theor. Comput. Sci. **754**, 16–34 (2019)
28. Izumi, T., Otachi, Y.: Sublinear-space lexicographic depth-first search for bounded treewidth graphs and planar graphs. In: Proceedings of 47th International Colloquium on Automata, Languages, and Programming, (ICALP 2020). LIPIcs, vol. 168, pp. 67:1–67:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2020). <https://doi.org/10.4230/LIPIcs.ICALP.2020.67>
29. Kammer, F., Kratsch, D., Laudahn, M.: Space-efficient biconnected components and recognition of outerplanar graphs. Algorithmica **81**(3), 1180–1204 (2019). <https://doi.org/10.1007/s00453-018-0464-z>
30. Kammer, F., Sajenko, A.: Simple  $2^f$ -color choice dictionaries. In: Proceedings of 29th International Symposium on Algorithms and Computation (ISAAC 2018). LIPIcs, vol. 123, pp. 66:1–66:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2018). <https://doi.org/10.4230/LIPIcs.ISAAC.2018.66>
31. Kammer, F., Sajenko, A.: Space efficient (graph) algorithms. <https://github.com/thm-mni-ii/sea> (2018)
32. Kammer, F., Sajenko, A.: Sorting and ranking of self-delimiting numbers with applications to tree isomorphism (2020)
33. Katoh, T., Goto, K.: In-place initializable arrays. CoRR [arXiv:1709.08900](https://arxiv.org/abs/1709.08900) (2017)
34. Lagergren, J.: Efficient parallel algorithms for graphs of bounded tree-width. J. Algorithms **20**(1), 20–44 (1996). <https://doi.org/10.1006/jagm.1996.0002>
35. Munro, J.I., Paterson, M.S.: Selection and sorting with limited storage. Theor. Comput. Sci. **12**(3), 315–323 (1980)
36. Munro, J.I., Raman, V.: Succinct representation of balanced parentheses, static trees and planar graphs. In: Proceedings of 38th Annual Symposium on Foundations of Computer Science (FOCS 1997), pp. 118–126. IEEE Computer Society (1997). <https://doi.org/10.1109/SFCS.1997.646100>

37. Reed, B.A.: Finding approximate separators and computing tree width quickly. In: Proceedings of 24th Annual ACM Symposium on Theory of Computing (STOC 1992), pp. 221–228. ACM (1992). <https://doi.org/10.1145/129712.129734>
38. Robertson, N., Seymour, P.: Graph minors xiii. The disjoint paths problem. *J. Comb. Theory Ser. B* **63**(1), 65–110 (1995). <https://doi.org/10.1006/jctb.1995.1006>
39. Robertson, N., Seymour, P.D.: Graph minors. II. Algorithmic aspects of tree-width. *J. Algorithms* **7**(3), 309–322 (1986). [https://doi.org/10.1016/0196-6774\(86\)90023-4](https://doi.org/10.1016/0196-6774(86)90023-4)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.