

List Colouring Trees in Logarithmic Space

Hans L. Bodlaender  

Utrecht University, The Netherlands

Carla Groenland  

Utrecht University, The Netherlands

Hugo Jacob  

ENS Paris-Saclay, France

Abstract

We show that LIST COLOURING can be solved on n -vertex trees by a deterministic Turing machine using $O(\log n)$ bits on the worktape. Given an n -vertex graph $G = (V, E)$ and a list $L(v) \subseteq \{1, \dots, n\}$ of available colours for each $v \in V$, a list colouring for G is a proper colouring c such that $c(v) \in L(v)$ for all v .

2012 ACM Subject Classification Theory of computation \rightarrow Graph algorithms analysis; Mathematics of computing \rightarrow Trees; Theory of computation \rightarrow Parameterized complexity and exact algorithms

Keywords and phrases List colouring, trees, space complexity, logspace, graph algorithms, tree-partition-width

Digital Object Identifier 10.4230/LIPIcs.ESA.2022.24

Related Version *Full Version:* <https://arxiv.org/abs/2206.09750>

Funding *Carla Groenland:* Supported by the project CRACKNP that has received funding from the European Research Council (grant agreement No 853234).

Acknowledgements We thank Marcin Pilipczuk and Michał Pilipczuk for discussions.

1 Introduction

Various applications can be modelled as an instance of LIST COLOURING, e.g., the vertices may correspond to communication units, with lists giving the possible frequencies or channels that a vertex may choose from as colours and edges showing which units would interfere if they are assigned the same colour [17, 24].

Given a graph $G = (V, E)$ and given a list $L(v)$ of colours for each vertex $v \in V$, an L -colouring c is a proper colouring (that is, $c(u) \neq c(v)$ when $uv \in E$) mapping every vertex v to a colour in the list $L(v)$. This gives rise to the following computational problem.

LIST COLOURING

Input: A graph $G = (V, E)$ with a list $L(v) \subseteq \{1, \dots, n\}$ of available colours for each $v \in V$.

Question: Is there an L -colouring for G ?

LIST COLOURING is computationally hard. It is NP-complete on cographs [19] and on planar bipartite graphs, even when all lists are of size at most 3 [18]. The problem remains hard when parameterised by “tree-like” width measures: it was first shown to be $W[1]$ -hard parameterised by treewidth in 2011 by [16] and recently shown to be XNLP-hard implying $W[t]$ -hardness for all t by [4]. On the other hand, on n -vertex trees the problem can be solved in time linear in n (using hashing) [19], but this algorithm may use $\Omega(n)$ space.



© Hans L. Bodlaender, Carla Groenland, and Hugo Jacob;
licensed under Creative Commons License CC-BY 4.0

30th Annual European Symposium on Algorithms (ESA 2022).

Editors: Shiri Chechik, Gonzalo Navarro, Eva Rotenberg, and Grzegorz Herman; Article No. 24; pp. 24:1–24:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In this paper, we study the auxiliary space requirements of LIST COLOURING on trees in terms of the number of vertices n of the tree. We assume that the vertices of T have been numbered $1, \dots, n$, which gives a natural order on them, and that, given vertices v, v' in T and $i, i' \in \{1, \dots, n\}$, it can be checked in $O(\log n)$ space whether the i th colour in $L(v)$ equals the i' th colour in $L(v')$. As is usual for the complexity class L (logspace), we measure the space requirements in terms of the number of bits on the work tape of a deterministic Turing machine, where the description of the tree and the lists are written on a (read-only) input tape. In particular, the number of bits on the input tape is allowed to be much larger.

Since n -vertex trees have pathwidth $O(\log n)$, our problem can be solved non-deterministically using $O(\log^2 n)$ bits on the work tape (see Proposition 3). However, doing better than this is surprisingly challenging, even in the non-deterministic case! Our main result is as follows.

► **Theorem 1.** LIST COLOURING for trees is in L.

Our initial interest in the space complexity of LIST COLOURING on trees arose from a recent result showing that LIST COLOURING parameterised by pathwidth is XNLP-complete [4]. XNLP is the class of problems on an input of size n with parameter k , which can be solved by a non-deterministic Turing machine in $f(k)n^{O(1)}$ time and $f(k)\log n$ space for some computable f . Since the treewidth of a graph is upper bounded by the pathwidth, LIST COLOURING is also XNLP-hard parameterised by treewidth. This is conjectured¹ to imply that there is a constant k^* for which any deterministic Turing machine needs $\omega(\log n)$ space in order to solve LIST COLOURING for n -vertex graphs of treewidth k^* ; this work shows that $k^* > 1$. It seems likely that LIST COLOURING parameterised by treewidth is not in XNLP, and we conjecture that it is complete for a parameterised analogue of NAuxPDA (also known as SAC) from [1, 23]. Considering such classes which (also) have space requirements (complexity classes such as XL, XNL and XNLP [4, 5, 7, 15]) has proven successful in classifying the complexity of parameterised problems which are not known to be complete for any classes that only consider time requirements. Since some of such classes are very naturally modelled by instances of LIST COLOURING, we believe the complexity class of LIST COLOURING on trees could be of theoretical interest as well.

Another motivation for studying space requirements comes from practice, since memory can be much more of a bottleneck than processing time (e.g. for dynamic programming approaches). This motivates the development of techniques to reduce the space complexity. Although many techniques have been established to provide algorithms which are efficient with respect to time, fewer techniques are known to improve the space complexity. Notable exceptions include the logspace analogue of Bodlaender's and Courcelle's theorem [14] which allows one to check any monadic second-order formula on graphs of bounded treewidth in logspace (which in particular allows one to test membership in any minor-closed family) and Reingold's [25] work on undirected connectivity. Reachability and isomorphism questions have also been well-studied on restricted graph classes, e.g. [8, 21, 22]. Another interesting piece of related work [13] shows that for each graph H , LIST H -COLOURING is either in L or NL-hard. We remark that LIST H -COLOURING on trees (for fixed H) is easily seen to be solvable in logspace using an analogue of Proposition 3 or the logspace Courcelle's theorem [14]. The difficulty in our case comes from the fact that the sizes of the lists are unbounded.

¹ The conjecture (see [23, Conjecture 2.1] or [4, Conjecture 5.1]) states that there is no deterministic algorithm for an XNLP-hard problem that runs in XP time and FPT space. If for each k , there exists a constant $f(k)$ such that LIST COLOURING can be solved in space $f(k)\log n$ on n -vertex graphs of treewidth k , then this would in particular yield an algorithm running in $n^{f(k)}$ time and $f(k)n^{O(1)}$ space.

We also generalise our algorithm to graphs of bounded tree-partition-width (also called strong treewidth).

► **Corollary 2.** *There is a deterministic $O(k \log k \log n)$ space algorithm for LIST COLOURING on n -vertex graphs with a given tree-partition of width k .*

The algorithm of Corollary 2 does not run in FPT time. We also include a simple proof that LIST COLOURING is $W[1]$ -hard when parameterised by tree-partition-width, which shows that it is unlikely that there exists an algorithm running in FPT time.

The algorithm of Theorem 1 is highly non-trivial and requires several conceptual ideas that we have attempted to separate out by first explaining some key ideas and an easier deterministic algorithm that uses $O(\log^2 n)$ space in Section 3. We assume our given tree to be rooted and transverse it by first “recursing” on children whose subtrees are not the largest. This bounds the number of such recursions by $O(\log n)$, and so we can “spend” $O(\log n)$ space per recursion. When we move to the heaviest subtree, we have either already rejected, or may forget entirely about the colour of v , or found a single colour that “works” for the non-heavy subtrees of v (“criticality”). Along a heavy branch, we always keep at most two such colours “per recursion depth” (the colour of v , and possibly one of its parent; once we move to the child of v we may forget the colour of the parent of v).

There are two further main ideas that remove the additional $\log n$ -factor. Most importantly, when we “move” from a vertex v to one of its children u , we will let the amount of storage allocated for storing v and its colour depend on the size of the subtree of u : the larger the subtree, the less space we allow. In the extreme case in which the size of the subtree of u is linear in the size of the subtree of v , we allow only a *constant* number of bits. At specific points during the algorithm, when more space is available temporarily, we use this stored information to recompute the vertex v and its colour. Suppose that v has d children u_1, \dots, u_d . We define the list $L_j(v)$ as the colours c such that when we give v colour c , we can extend the colouring to the subtrees of u_1, \dots, u_j . If $|L_j(v)| > d - j$, then v is “non-critical”: we will always be able to assign it a colour after colouring the subtrees of u_{j+1}, \dots, u_d . This allows us to maintain that $|L_j(v)| \leq d - j$ and so the number of bits required to store the position of c in $L_j(v)$ decreases as j increases.

We need to be a bit more clever when we define the lists. The second main idea is to distribute the children of v into about $\log \log(n/2)$ brackets, where n denotes the number of vertices in the subtree below v . The distribution is done based on how much smaller the subtree of the child is compared to n . We allocate a specific number of bits per bracket: to brackets which allow bigger subtrees, we allocate less space. When “processing” brackets of smaller subtrees, we may need to store information as well for the brackets of bigger subtrees, but vice versa is not allowed. We choose the bracket sizes so that if we store information for the first j brackets, this “fits” in the space allocated to the $(j + 1)$ th bracket. In the end, the final algorithm is rather subtle and requires a careful analysis.

We outline some relevant definitions and background in Section 2. We give the simpler $O(\log^2 n)$ space algorithm in Section 3 and discuss the logspace algorithm in Section 4. We prove our results concerning tree-partition-width in Section 5 and point to some directions for future work in Section 6. Some technical details can be found in the full version [3].

2 Preliminaries

All logarithms in this paper have base 2. Let T be a rooted tree and $v \in V(T)$. We write T_v for the subtree rooted at v and $T - v$ for the forest obtained by removing v and all edges incident to v . We refer the reader to textbooks for basic notions in graph theory [9] and (parameterised) complexity [2, 12].

2.1 Simple logspace computations on trees

We repeatedly use the fact that simple computations can be done on a rooted tree using logarithmic space, such as counting the number of vertices in a subtree. We include a brief sketch below and refer to [22] for further details.

We first explain how to traverse a tree in logspace. Record the index of the current vertex and create states **down**, **next** and **up**. We start on the root with state **down**. When in state **down**, we go to the first child while remaining in the state **down**. If there is no child, we change the state to **next**. When in state **next**, we go to the next sibling if it exists and change state to **down**, or (if there is no next sibling) we change state to **up**. When in state **up**, we simultaneously go to the parent and change the state to **next**. We stop when reaching the root with state **up**. By keeping track of the number of vertices discovered, we can use the same technique to count the number of vertices in a subtree. This can then be used to compute the child with maximum subtree size and to enumerate children ordered by their subtree size. If the input tree is not rooted, we may use the indices of the vertices to root the tree in a deterministic way.

2.2 Graph width measures

Let $G = (V, E)$ be a graph. A tuple $(T, \{X_t\}_{t \in V(T)})$ is a *tree decomposition* for G if T is a tree, for $t \in V(T)$, $X_t \subseteq V$ is the *bag* of t , for each edge $uv \in E$ there is a bag such that $\{u, v\} \subseteq X_t$, and for each $u \in V$, the bags containing u form a nonempty subtree of T . If T is a path, this defines a *path decomposition*.

The width of such a decomposition is $\max_{t \in V(T)} |X_t| - 1$. The *treewidth* (resp. *pathwidth*) of G is the minimum possible width of a tree decomposition (resp. path decomposition) of G .

A *nice* path decomposition has empty bags on endpoints of the path and two consecutive bags differ by at most one vertex. Hence, either a vertex is *introduced*, or a vertex is *forgotten*.

Let G be a graph, let T be a tree, and, for all $t \in V(T)$, let X_t be a non-empty set so that $(X_t)_{t \in V(T)}$ partitions $V(G)$. The pair $(T, (X_t)_{t \in V(T)})$ is a *tree-partition* of G if, for every edge $vv' \in E(G)$, either v and v' are part of the same bag, or $v \in X_t$ and $v' \in X_{t'}$ for $tt' \in E(T)$. The width of the partition is $\max_{t \in V(T)} |X_t|$. The *tree-partition-width* (also known as strong treewidth) of G is the minimum width of all tree-partitions of G . It was introduced by Seese [26] and can be characterised by forbidden topological minors [11]. Tree-partition-width is comparable to treewidth on graphs of maximum degree Δ [10, 27]: $\text{tw} + 1 \leq 2 \text{tpw} \leq O(\Delta \text{tw})$. However, it is incomparable to treedepth, pathwidth and treewidth for general graphs.

The *treedepth* of a graph is the minimum height of a forest F with the property that every edge of G connects a pair of nodes that have an ancestor-descendant relationship to each other in F .

3 Warm up: first ideas and a simpler algorithm

3.1 Storing colours via their position in the list

It is not too difficult to obtain a non-deterministic algorithm that uses $O(\log^2 n)$ space.

► **Proposition 3.** *LIST COLOURING can be solved non-deterministically using $O(\log n \log \Delta)$ space on n -vertex trees of maximum degree Δ .*

The proposition follows from the following two lemmas.

► **Lemma 4.** *LIST COLOURING can be solved non-deterministically using $O(k \log \Delta + \log n)$ space for an n -vertex graph G of maximum degree Δ if we can deterministically compute a path decomposition for G of width k in $O(\log n)$ space.*

A deterministic logspace algorithm for computing an optimal path decomposition exists for all graphs of bounded pathwidth [20], but this does not apply directly to trees (since their pathwidth may grow with n).

► **Lemma 5.** *If T is an n -vertex tree, we can deterministically construct a nice path decomposition of width $O(\log n)$ using $O(\log n)$ space.*

We remark that Δ may be replaced by a bound on the list sizes in Proposition 3 and Lemma 4. The main observation in the proof of Lemma 4 is that for a vertex v , we only need to consider the first $d(v) + 1$ colours from its list so that we can store the *position* of the colour rather than the colour itself. Note that we cannot keep the path decomposition in memory, but rather recompute it whenever any information is needed. We keep in memory only the current position in the path decomposition and the list positions of the colours we assigned for vertices in the previous bag.

3.2 Heavy children, recursive analysis and criticality

Suppose that we are given an instance (T, L) of LIST COLOURING. We fix a root v^* of T in an arbitrary but deterministic fashion, for example the first vertex in the natural order on the vertices. Let $v \in V(T)$. We see v as a descendant and ancestor of itself. We write T_v for the subtree with root v .

► **Definition 6** (Heavy). *A child u of a vertex v in a rooted tree T is called heavy if $|V(T_u)| \geq |V(T_{u'})|$ for all children u' of v , with strict inequality whenever $u' < u$ in the natural order on V .*

Each vertex has at most one heavy child. We also record the following nice property.

► **Observation 7.** *If u is a child of v which is non-heavy, then $|V(T_u)| \leq (|V(T_v)| - 1)/2$.*

An obvious recursive approach is to loop over the possible colour $c \in L(r)$ for the root r and then to recursively check for all children v of r whether a list colouring can be extended to the subtree T_v (while not giving v the colour c). We wish to prove a space upper bound of the form $S(n) = f(n) \log n$ on the number of bits of storage required for trees on n vertices (for some non-decreasing function f). We compute

$$S(n/2) = \log(n/2)f(n/2) \leq \log(n/2)f(n) = \log n f(n) - \log 2 f(n) \leq S(n) - f(n). \quad (1)$$

This shows that while performing a recursive call on some subtree T_v with $|V(T_v)| = n/2$, we may keep an additional $f(n)$ bits in memory (on top of the space required in the recursive call). In particular, we can store the colour c using $O(\log n)$ bits when $f(n) = \Theta(\log n)$, but can only keep a *constant* number of bits for such recursions when proving Theorem 1.

We next explain how we can ensure that we only need to consider recursions done on non-heavy children. Suppose v has non-heavy children v_1, \dots, v_k and heavy child u . We will write $G_v = T_v - T_u$. Suppose the parent v' of v needs to be assigned colour c' . One of the following must be true.

- There is no colouring of G_v which avoids colour c' for v . In this case, we can reject.
- There is a unique colour $c \neq c'$ that can be assigned to v in a list colouring of G_v . We say v is *critical* and places the colour constraint on u that it cannot receive colour c .

- There are two possible colours unequal to c' that we can give v in G_v . We then say v is *non-critical*: it can be coloured for each colour we might wish to assign its heavy child u . One example of non-criticality is if the list of available colours $|L(v)| \geq k + 2$. In this case, T_v is list colourable if and only if $T_{v_1}, \dots, T_{v_k}, T_u$ are all list colourable (using lists L).

3.3 A deterministic algorithm using $O(\log^2 n)$ space and polynomial time

We define a procedure $\text{SOLVE}(v, p)$ which given a vertex v and number $p \in \{0, 1, \dots, n\}$, determines whether the subtree T_v rooted on v can be list coloured; for $p \geq 1$ there is an additional constraint that v cannot receive the p th colour in $L(v)$; for $p = 0$ the vertex v may receive any colour.

Suppose we call $\text{SOLVE}(r, p)$ for some $r \in V(T)$. Let v_1, \dots, v_k denote the non-heavy children of r and u the heavy child. The algorithm works as follows.

1. For $i = 1, \dots, k$, we recursively verify that T_{v_i} can be list coloured ($\text{SOLVE}(v_i, 0)$); we reject if any of these rejects.
2. If $|L(r)| \geq d(r) + 1 = k + 2$, then we are “non-critical”: we free up our memory (removing also (r, p)) and recursively verify whether T_u can be list coloured by calling $\text{SOLVE}(u, 0)$.
3. From now on, $|L(r)| \leq k + 1$. We check that there is some $p_1 \in \{1, \dots, |L(r)|\}$ for which we can assign v the p_1 th colour in its list, and extend to a list colouring of $T \setminus T_u$. This involves recursive calls $\text{SOLVE}(v_i, p_{i,1})$ where $p_{i,1}$ places the appropriate constraint² on v_i . If no such p_1 exists, we reject.
4. Next, we check whether r is “non-critical”, that is, whether there is some $p_2 \neq p_1$ for which there is a list colouring of $T \setminus T_u$ in which r receives the p_2 th colour from its list. If such p_2 exists, we free up our memory and recursively verify whether T_u can be list coloured. If p_2 does not exist, then we know that r must get colour p_1 . If $p_1 = p$, we reject. Otherwise, we free up our memory and run $\text{SOLVE}(u, p'_1)$, where p'_1 is either 0 or the position in the list of u of the p_1 th colour of $L(r)$.

We give a brief sketch of the space complexity; more precise arguments including also pseudocode and the time analysis of this algorithm are given in the full version [3].

For the sake of analysis, suppose we keep a counter \mathbf{r} that keeps track of the “recursion depth”. We increase this by one each time we do a call on a non-heavy child (decreasing it again once it finishes), but do not adjust it for calls on a heavy child.

Suppose \mathbf{r} has been increased due to a sequence of recursive calls on $(v_1, 0), (v_2, p_2), \dots, (v_\ell, p_\ell)$, with v_1 the root of the tree and v_{i+1} a non-heavy child of some heavy descendant of v_i for all $i \in [\ell - 1]$. Then $|V(T_{v_\ell})| \leq \frac{1}{2}|V(T_{v_{\ell-1}})| \leq \dots \leq \frac{1}{2^{\ell-1}}|V(T)|$. In particular, \mathbf{r} is always bounded by $\log n$.

Crucially, if a call $\text{SOLVE}(u, p')$ is made on the heavy descendant u of some v_i , the only information we need to store relating to the part of the tree “between v_i and u ” is p' . Therefore, if we distribute our work tape into $\lceil \log n \rceil$ parts where the i th part will be used whenever \mathbf{r} takes the value i , then each part only needs to use $10 \log n$ bits, giving a total space complexity of $O(\log^2 n)$.

² Let c_1 be the p_1 th colour in $L(r)$. If $c_1 \notin L(v_i)$, let $p_{i,1} = 0$. Otherwise, let $p_{i,1}$ be the position of c_1 in $L(v_i)$.

4 Proof of Theorem 1

In this section, we describe our $O(\log n)$ space algorithm. This also uses the ideas of using positions in a list (rather than the colours themselves), criticality and starting with the non-heavy children described in the previous section. However, we need to take the idea of first processing “less heavy” children even further.

The main idea is to store the colour c that we are trying for a vertex v using the position p_j of c in some list $L_j(v)$, and to reduce the size of the list (and therefore the storage requirement of p_j) before we process “heavier” children of v . There are two key elements:

- We can recompute c from p_j in $O(\log n)$ space. This is useful, since we can do a recursive call while only having p_j (rather than c) as overhead, discover some information, and then recompute c only at a point where we have a lot of memory available to us again.
- The space used for p_j will depend on the size of the tree that we process. It is too expensive to consider all sizes separately, and therefore we will “bracket” the sizes. Subtrees whose size falls within the same bracket are processed in arbitrary order. For example, we put all trees of size $O(\sqrt{n})$ in a single bracket: these can be processed while using $O(\log n)$ bits of information about c (which is trivially possible). At some point we reach brackets for which the subtrees have size linear in n , say of size at least $\frac{1}{4}n$. Then, we may only keep a *constant* number of bits of information about c . Intuitively, this is possible because at most four children of r can have a subtree of size $\geq \frac{1}{4}n$, so the “remaining degree” of v is small. In particular, if more than six colours for v work for all smaller subtrees, then v is “non-critical”.

We first explain our brackets in Section 4.1. We then explain how we may point to a colour using less memory in Section 4.2 and how we keep track of vertices using less memory in Section 4.3. We then sketch the proof of Theorem 1 by outlining the algorithm and its space analysis in Section 4.4.

4.1 Brackets

Recall that we fixed a root for T in an arbitrary but deterministic fashion.

Let $v \in V(T)$ and u the heavy child of v . Let $G_v = T_v - T_u$ and $n_v = |V(T_v)|$. Each subtree T' of $G_v - v$ is rooted in a non-heavy child of v and will be associated to a bracket based on $|V(T')|$. By Observation 7, $1 \leq |V(T')| \leq (n_v - 1)/2$. Let $M_v = \lceil \log \log(n_v/2) \rceil$. The brackets are given by the sets of integers in the intervals

$$[1, n_v/2^{2^{M_v-1}}), [n_v/2^{2^{M_v-1}}, n_v/2^{2^{M_v-2}}), \dots, [n_v/256, n_v/16), [n_v/16, n_v/4), [n_v/4, n_v/2).$$

There are M_v brackets: $[n_v/2^{2^j}, n_v/2^{2^{j-1}})$ is the j th bracket for $j \in \{1, \dots, M_v - 1\}$ and $[1, n_v/2^{2^{M_v-1}})$ is the M_v th bracket. Note that $n_v/2^{2^{M_v-1}} = O(\sqrt{n_v})$. This implies that while doing a recursive call on a tree in the M_v th bracket, we are happy to keep an additional $O(\log n_v)$ bits in memory.

We aim to show that for some universal constant C , when doing a recursive call on a tree in the j th bracket, we can save all counters relevant to the current call using at most $C2^j$ bits (which depending on the value of j , could be $\Omega(\log n)$). In our analysis, we save the counters in a new read-only part of the work tape. The recursive call cannot alter this (and will have to work with less space on the work space). We can then use our saved state to continue with our calculations once the recursive call finishes.

We short-cut $M = M_v$ for legibility; the dependence of M on v is only needed to ensure that we do not start storing counters for lots of empty brackets when n_v is much smaller than n , and can be mostly ignored.

4.2 The information p_j stored about a colour c

Let $v \in V(T)$ with heavy child u and $c \in L(v)$. Recall that $G_v = T_v - T_u$. We will loop over $j = M, \dots, 0$ and consider subtrees of G_v whose size falls in the j th bracket in an arbitrary, but deterministic order (e.g. using the natural order on their roots). When j decreases, we will perform recursions on larger subtrees of $G_v - v$ and can therefore keep less information about c . We first define “implicit” lists.

- Set $L_M(v) = L(v)$.
- For $j \in [0, M - 1]$, let $L_j(v)$ be the set of colours $\alpha \in L(v)$ such that all subtrees T' of $G_v - v$ with $|V(T')| < n/2^{2^j}$ can be coloured without giving the colour α to the root of T' (so that v may receive colour α according to those subtrees).

Note that $L_j(v) = L(v)$ if there are no subtrees T' with $|V(T')| < n/2^{2^j}$. Since the subtrees associated to brackets $1, \dots, j$ have size at least $n/2^{2^j}$, there can be at most 2^{2^j} of them. If $|L_j(v)| \geq 2^{2^j} + 3$ and all subtrees T' of $G_v - v$ can be coloured, then v is “non-critical”: after the parent and heavy child of v have been coloured, the colouring can always be extended to v and the rest of G_v .

Suppose we are testing if we can give colour c to v . If $c \notin L_j(v)$, then we may reject: c is not a good colour for one of the subtrees. We define $p_j = p_j(c, v)$ as the integer $x \in \{1, \dots, |L_j(v)|\}$ such that the x th element in $L_j(v)$ equals c . In particular, p_M is the position of the colour c in the list $L_M(v) = L(v)$ and p_0 gives the position of c in the list of colours for which all subtrees of G_v allow v to receive c . For $j < M$, we will reserve at least $\log(2^{2^j} + 3)$ bits for p_j . This is possible, because we can maintain that $|L_j(v)| \leq 2^{2^j} + 2$ by going into a “non-critical subroutine” if we discover the list is larger.

4.3 Position of the current vertex

Next, we describe how to obtain efficient descriptions of the vertices in the tree. When performing recursions, we find it convenient to store information using which we can retrieve the “current vertex” of the parent call. Therefore, we require small descriptions for such vertices if the call did not make much “progress”.

For any $v \in V(T)$, define a sequence $h(v, 1), h(v, 2), \dots$ of heavy descendants as follows. Let $h(v, 1) = v$. Having defined $h(v, i)$ for some $i \geq 1$, if this is not a leaf, we let $h(v, i + 1)$ be the heavy child of $h(v, i)$. Note that given the vertex $h(v, i)$, we can find the vertex $h(v, i + 1)$ (or conclude it does not exist) in $O(\log n)$ space. We give a deterministic way of determining a bit string $\text{pos}(v, i)$ that represents $h(v, i)$, where the size of the bit string will depend on the “progress” made at the vertex $h(v, i)$ that it represents. This “progress” is measured by the size t_i of the largest subtree T' of a non-heavy child of $h(v, i)$, that is, T' is the largest component of $T - h(v, i)$ which does not contain $h(v, j)$ for $j \neq i$. We define $\text{pos}(v, i)$ as follows.

- Let j be given such that $t_i \in [n_v/2^{2^j}, n_v/2^{2^{j-1}})$. Start $\text{pos}(v, i)$ with j zeros, followed by a 1.
- There are at most 2^{2^j} values of i for which $t_i \geq 2^{2^j}$. We add a bit string of length 2^j to $\text{pos}(v, i)$, e.g. the value x for which $h(v, i)$ is the x th among $h(v, 1), \dots, h(v, a)$ with $t_i \in [n_v/2^{2^j}, n_v/2^{2^{j-1}})$.

Note that, given v , we can compute $\text{pos}(v, i)$ from $h(v, i)$ and $h(v, i)$ from $\text{pos}(v, i)$ using $O(\log n)$ space. If we do a recursive call, it will be on a non-heavy child u of some $h(v, i)$. By definition, $|V(T_u)| \leq t_i$ and $\text{pos}(v, i)$ depends on t_i in a way that we are able to keep it in memory while doing the recursive call.

We use the encoding $(\text{pos}(v^*, i), j, \ell)$ for the ℓ th child u of $h(v^*, i)$ whose subtree has a size that falls in the j th bracket. We can also attach another such encoding, e.g.

$$((\text{pos}(v^*, i), j, \ell), (\text{pos}(u, i'), j', \ell')),$$

to keep track of u and the ℓ' th child v of $h(u, i')$ whose subtree has a size that falls in the j' th bracket. We can retrieve v from the encoding above in $O(\log n)$ space and therefore can retrieve it whenever we have such space available to us.

4.4 Description of the algorithm

During the algorithm, the work tape will always start with the following.

- \mathbf{r} : the recursion depth r written in unary. At the start, $r = 0$.
- $\mathbf{pos} = \text{pos}_0 | \dots | \text{pos}_r$: encodes vertices as described Section 4.3. At the start, this is empty and points at the root v^* of the tree on the input tape.
- $\mathbf{p} = \mathbf{p}_0 | \dots | \mathbf{p}_r$: encodes colour restriction information for the vertices encoded by \mathbf{pos} . At the start, this is empty and no restrictions are given. We maintain throughout the algorithm that \mathbf{p}_i gives us colour restrictions for the vertex v pointed at by pos_i . The restriction can either be “no restrictions” or “avoid c' ”; in the latter case \mathbf{p}_i contains a tuple (j, p'_j) with p'_j the position of c' in $L_j(v')$, where v' is the parent of v .
- $\mathbf{aux} = \mathbf{aux}_0 | \dots | \mathbf{aux}_r$: further auxiliary information for parent calls that may not be overwritten.

We define a procedure **process**. While the value of \mathbf{r} equals r , the bits allocated to $\text{pos}_i, \mathbf{p}_i, \mathbf{aux}_i$ for $i < r$ will be seen as part of the read-only input-tape. In particular, the algorithm will not make any changes to $\text{pos}_i, \mathbf{p}_i$ or \mathbf{aux}_i for any $i < r$, but may change the values for $i = r$.

We will only increase \mathbf{r} when we do a recursive call. If during the run of the algorithm $\mathbf{r} = r$ and a recursive call is placed, then we increase \mathbf{r} to $r + 1$ and return to the start of our instructions. However, since \mathbf{r} has increased it will now see a “different input tape”. When the call finished, we will decrease \mathbf{r} back to r and wipe everything from the work space except for $\mathbf{r}, \text{pos}_i, \mathbf{p}_i, \mathbf{aux}_i$ for $i \leq r$, and the answer of the recursive call (0 or 1). We then use \mathbf{aux}_i to reset our work space and continue our calculations.

We will ensure that \mathbf{r} is always upper bounded by $\log n$. Indeed, the vertex v_r encoded by pos_r will always be a non-heavy child of a descendant of the vertex encoded by pos_{r-1} .

While $\mathbf{r} = r$, the algorithm is currently doing calculations to determine whether the vertex v pointed at by pos_{r-1} (v^* for $r = 0$) has the property that T_v can be list coloured, while respecting the colour restrictions encoded by \mathbf{p}_{r-1} (none if $r = 0$). Recall that rather than writing down v explicitly, we use a special encoding from which we can recompute v whenever we have $C \log n$ space available on the work tape (for some universal constant C). Similarly, \mathbf{p}_{r-1} may give a position p'_j in $L_j(v')$ for some $j < M$ (for v' the parent of v), and we may need to use our current work tape to recompute the corresponding position p'_M of the colour in $L(v')$, so that we can access the colour from the input tape. This part will make the whole analysis significantly more technical.

We define an algorithm which we call **process** as follows. A detailed outline is given in the full version [3], whereas an informal description of the steps is given below.

0. Let v be the vertex pointed at by pos_r . We maintain that at most one colour c' has been encoded that v must avoid. Handle the case in which v is a leaf. If not, it has some heavy child h . We go to 1, which will eventually lead us to one of the following (recall that $G_v = T_v - T_h$):

(**rej**) There is no list colouring of G_v avoiding c' for v . We return **false**.

- (nc) The vertex v can get two colours (unequal to c') in list colourings of G_v . In this case, we say v is non-critical. We update pos_r to h , set p_r to “none” and repeat from 0.
- (cr) There is a unique colour $c \neq c'$ that works. Then $\{c\} \subseteq L_0(v) \subseteq \{c, c'\}$ and so can represent $p_0 = p_0(c, v)$ with a single bit. We update pos_r to h , update p_r to $(0, p_0)$ and repeat from 0.
1. We check that all subtrees can be coloured if we do not have any colour restrictions (necessary for the non-critical subroutine). This involves recursive calls on **process** where we have no colour constraint on the root of the subtree.
 2. We verify that $L_0(v)$ is non-empty. We iteratively try to compute p_M from $p_0 = 1$ via p_1, \dots, p_{M-1} . Starting from $p_0 = 1$ and $j = 0$, we compute p_{j+1} from p_j as follows.
 - (i) Initialise $\text{curr}_j = 1$. This represents a position in $L_j(v)$ (giving the number of “successes”).
Initialise $\text{prev}_j = 1$. This represents a position in $L_{j+1}(v)$ (giving the number of “tries”).
 - (ii) We check whether the prev_j th colour of $L_{j+1}(v)$ works for the trees in the j th bracket. This involves a recursive call on **process** for each tree T' in the j th bracket, putting $(j + 1, \text{prev}_j)$ as the colour constraint for the root of T' . (The colour restriction gives a position in $L_{j+1}(v)$; we do not store the corresponding colour and rather will recompute it in the recursive call!)
 - (iii) If one of those runs fails, we increase prev_j ; if this is now $> 2^{2^{j+1}} + 3$, then this implies a lower bound on $|L_{j+1}(v)|$ which allows us to move to (nc).
If $\text{curr}_j < p_j$ we increase both curr_j and prev_j .
Once $\text{curr}_j = p_j$, we have successfully computed $p_{j+1} = \text{prev}_j$ and continue to compute p_{j+2} if $j + 1 < M$. Otherwise we repeat from (ii).

Once we have p_M , we can access the corresponding colour from the input tape by looking at the p_M th colour of $L(v)$. If $p_M > |L(v)|$, we go to (rej).

3. We verified $|L_0(v)| \geq 1$. We establish whether $|L_0(v)| \geq 3$ in a similar manner. If so, we can go to (nc); else we need to start considering the colour constraints of the parent v' of v . Note that we can use aux_r to store auxiliaries such as $\alpha = |L_0(v)| \in \{1, 2\}$.

It remains to explain how we check whether the first or second colour from $L_0(v)$ satisfies the colour constraint from v' . Suppose a colour c' has been encoded via the position $p'_{j'}$ of c' in $L_j(v')$ for some $j' \in [0, M']$ (where $M' = M_{v'}$).

We can recompute $p'_{M'}$ from $p'_{j'}$ in a similar manner to the above. However, once we store $p'_{M'}$, we may no longer be able to compute p_M from $p_0 = 1$ or $p_0 = 2$, since $p'_{M'}$ may take too much space³. Therefore, we first recompute $p_{j'}$ from p_0 and then simultaneously recompute $p_{j'+1}$ from $p_{j'}$ and $p'_{j'+1}$ from $p'_{j'}$ until we computed p_M and $p'_{M'}$. We then check whether the p_M th colour of $L(v)$ equals c' , the $p'_{M'}$ th colour of $L(v')$. (It may be that $M \neq M'$, meaning that we may finish one before the other.)

The computation of p'_{x+1} from p'_x for the parent v' of v is a bit more complicated if v is a non-heavy child of v' . In this case, v is in the $(j' - 1)$ th bracket of v' . The algorithm calls again on itself for subtrees in the x th bracket of v' , but now we see a resulting call to **process** as a *same-depth call* rather than a “recursive call”. The computations work the same way, but we do not adjust r and will *add* the current state from before the call in aux_r .

³ This is one of the issues that made this write-up more technical and involved than one might expect necessary at first sight; if we do a recursion on a child in the j th bracket of v or v' , then we are only allowed to keep $O(2^j)$ bits on top of the space used by the recursion. This means we cannot simply keep $p'_{M'}$ in memory if j is much smaller than M' .

By an exhaustive case analysis, the algorithm computes the right answer. It remains to argue that it terminates and runs in the correct space complexity.

We first further explain the same-depth calls on `process`. When we make such call, we do not adjust \mathbf{r} . When the call is made, \mathbf{pos}_r will encode a vertex v_1 which is a non-heavy child of v' . After the call, \mathbf{pos}_r will point to some other non-heavy child v_2 of v' . Importantly, the bracket of v_2 will always be higher than the bracket of v_1 , so that at most M such same-depth calls are made before changing our recursion level. Each same-depth call may stack on a number of auxiliary counters which we keep track of using \mathbf{aux}_r ; since for each j , there is at most one vertex from the j th bracket which may append something to this, it suffices to ensure a vertex from bracket j adds at most $C2^j$ bits. Indeed, this ensures that the size of \mathbf{aux}_r takes at most $C \sum_{j=1}^{j'} 2^j \leq C2^{j'+1} = O(2^{j'})$ bits once \mathbf{pos}_r encodes a vertex from bracket j' .

We now argue that it terminates. Each time we do a same depth call, the value of the bracket j will increase by at least one. We therefore may only do a finite number of same-depth calls in a row. When $\mathbf{r} = 0$, each time we reach (nc) or (cr), we move one step on the heavy path from the root to a leaf, so this will eventually terminate. A similar observation holds for $\mathbf{r} > 0$ once we fix $\mathbf{pos}_1 | \dots | \mathbf{pos}_{r-1}$: this points to some vertex v and \mathbf{pos}_r will initially point to a non-heavy descendent u or v , and then “travel down the heavy path” from u to a leaf.

We next consider the space used by the algorithm. Let $S(n)$ be the largest amount of bits used for storing $\mathbf{r}, \mathbf{pos}, \mathbf{p}$ and \mathbf{aux} during a run of `PROCESS` on an n -vertex tree. We can distribute our work space into two parts: $C_1 \log n$ space for temporary counters and for doing calculations such as computing the heavy child of a vertex, and $S(n)$ bits for storing $\mathbf{r}, \mathbf{pos}, \mathbf{p}$ and \mathbf{aux} . It suffices to prove that $S(n) \leq C_2 \log n$; this is the way we decided to formalise keeping track of the “overhead” caused by recursive calls.

Note that \mathbf{r} is bounded by $\log n$ if the input tree has n vertices: each time it increases, we moved to a non-heavy child whose subtree consists of at most $n/2$ vertices.

Suppose we call `process` with \mathbf{pos} pointing at some vertex v whose subtree has size n_v . We will show inductively that the number of bits used by \mathbf{pos}, \mathbf{p} and \mathbf{aux} is in $O(\log n_v)$ throughout this call (note that n_v may be much smaller than n , the number of vertices of the tree on the input tape). Whenever we do a recursive call, this will be done on a tree whose size is upper bounded by $n_v/2^{2^{j-1}}$ for some $1 \leq j \leq M = \lceil \log \log(n_v/2) \rceil$. Since by induction the recursive call requires only

$$S(n/2^{2^{j-1}}) = C_2 \log(n/2^{2^{j-1}}) \leq S(n) - C_2 2^{j-1}$$

additional bits, we will allow ourselves to add at most $C_2 2^{j-1}$ bits to $\mathbf{r}, \mathbf{pos}, \mathbf{p}$ and \mathbf{aux} before we do a recursive call that divides the number of vertices by at least $2^{2^{j-1}}$. The constant C_2 will be relatively small (for example, 1000 works).

Fix a value of j . From the definitions in Section 4.2 and 4.3, at the point that we do a recursion on a subtree whose size is upper bounded by $n_v/2^{2^{j-1}}$, \mathbf{pos}_r and \mathbf{p}_r store at most three integers (they are of the form $(\mathbf{pos}(v', i), j, \ell)$ and (x, p_x) respectively) that are bounded by $2^{2^j} + 3$. Therefore, these require at most $C' 2^{j-1}$ bits for some constant C' (e.g. $C' = 50$ works).

The worst case comes from \mathbf{aux}_r which may get stacked up on by the “same-depth calls”. There is at most one such same-depth call per $x \in \{1, \dots, j\}$. For such x , we add on a bounded number of counters (e.g. \mathbf{curr}_x and \mathbf{prev}_x) which can take at most $2^{2^{x+1}} + 3 \leq 2 \cdot 2^{2^{x+1}}$ values. Since $C \sum_{x=1}^j 2^x \leq 4C2^{j-1}$, \mathbf{aux}_r also never contains more than $O(2^{j-1})$ bits while doing a call on a tree whose size falls in bracket j .

5 Graphs of bounded tree-partition-width

5.1 Proof of Corollary 2

Here we sketch the proof of Corollary 2. Let $(T, (X_t)_{t \in V(T)})$ be a tree-partition of width k for an N -vertex graph G , where it will be convenient to write $n = |V(T)| \leq N$. We prove that there is an algorithm running in $O(k \log k \log n)$ space, for each k by induction on n .

We root T in the vertex of lowest index. We define the recursion depth, heavy children, the brackets, M and “position” for the vertices in T exactly as we did in Section 4.

Suppose pos points at some vertex $t \in V(T)$. We aim to use $O(k \log k \log(2^{2^j}))$ bits for the information p_j stored about the colouring c of the vertices in X_t while processing subtrees in bracket j . This is done as follows. For $t \in V(T)$, let $G_{t,j}$ be the graph induced on the vertices that are either in X_t or in X_s for s a non-heavy child of t in bracket j or above. For $v \in X_t$, we define $L_j(v)$ to be the list of colours $\alpha \in L(v)$ such that there is a list colouring of $G_{t,j}$ that assigns the colour α to v . There are at most 2^{2^j} subtrees of T associated to brackets $1, \dots, j$, and so these include at most $k2^{2^j}$ neighbours of v . Therefore, if $|L_j(v)| \geq k(2^{2^j} + 3)$, then G can be list coloured if and only if $G - v$ can be list coloured and we no longer need to keep track of the colour of v . We then establish v is non-critical. We use k bits to write for each vertex of X_t whether or not it has been established to be critical, and for those vertices that are critical, we use $\log(k(2^{2^j} + 3))$ bits per vertex to index a colour from $L_j(v)$.

We use $k + \log((3k)^k) = O(k \log k)$ bits of information in aux to keep track of the following.

- For each $v \in X_t$, whether or not v has been established to be critical. After processing t , the vertex $v \in X_t$ will be critical if there are at most $3k$ colours in $L(v)$ for which there exists an extension to G_t . Let $C_t \subseteq X_t$ denote the critical vertices.
- For each $p_0 \in \prod_{v \in C_t} L_0(v)$ (of which there are at most $(3k)^k$), a single bit which indicates whether or not the parent of t would allow the corresponding colouring.

While computing the information above, we still need the auxiliary information from the parent of t , but we can discard this by the point we start processing the heavy child of t .

We make two more small remarks:

- We need to redefine what we mean by “increasing” p_j for some j , since we now work with a tuple of list positions. We fix an arbitrary but deterministic way to do this, for example in lexicographical order using the natural orders on the vertices and colours.
- When we check whether the colouring c of X_t corresponding to $p_0 \in \prod_{v \in C_t} L_0(v)$ is allowed by the parent t' of t , we run over $p'_0 \in \prod_{v' \in C_{t'}} L_0(v')$, and as before we need to compute p_i, p'_i from p_{i-1}, p'_{i-1} iteratively until we obtain the colourings corresponding to p_M and p'_M . If those colourings are compatible, then we know that there is a list colouring of the graph “above t ” for which X_t is coloured “according to p_0 ”, and so we record in aux that p_0 is allowed.

The calculations in the space analysis work in the exact same way: we simply multiply everything by $Ck \log k$ (for a universal constant C).

5.2 W[1]-hardness

We give an easy reduction for the following result.

► **Theorem 8.** *LIST COLOURING parameterised by the width of a given tree-partition is W[1]-hard.*

Proof. We reduce from MULTICOLOURED CLIQUE.

Consider a MULTICOLOURED CLIQUE instance $G = (V, E), V_1, \dots, V_k$ with $k \geq 2$ colours. We denote by $\overline{G} = (V, \overline{E})$ the complement of G .

We now describe the construction of our instance graph H . We first add vertices v_1, \dots, v_k , with lists $L(v_i) = V_i$ for all $i \in [k]$. Then for each edge $e = uv \in \overline{E}$ we add a vertex x_{uv} with list $L(x_{uv}) = \{u, v\}$. Furthermore, for $\alpha \in \{u, v\}$, let i be such that $\alpha \in V_i$. We add the edge $x_{uv}v_i$.

The resulting graph has tree-partition-width at most k : we put v_1, \dots, v_k in the same bag which is placed at the centre of a star, and create a separate leaf bag containing x_{uv} for each $uv \in \overline{E}$.

▷ **Claim 9.** If there is a proper list colouring of H , then there is a multicoloured clique in G .

Proof. Suppose that H admits a list colouring. Let $a_i \in V_i = L(v_i)$ be the colour assigned to v_i for all $i \in [k]$. We will prove a_1, \dots, a_k forms a multicoloured clique in G .

Consider distinct $i, j \in [k]$ and suppose $a_i a_j$ is not an edge of G , that is, $a_i a_j \in \overline{E}$. Then there exists a vertex $x_{a_i a_j}$ adjacent to both v_i and v_j , but there is no way to properly colour it, a contradiction. So we must have $a_i a_j \in E$ as desired. ◁

▷ **Claim 10.** If there is a multicoloured clique in G , then there is a proper list colouring of H .

Proof. We denote by a_1, \dots, a_k the vertices of the multicoloured clique, where $a_i \in V_i$ for all i . We assign the colour a_i to vertex v_i . Consider now x_{uv} for some $uv \in \overline{E}$. Let i and j be given so that x_{uv} is adjacent to v_i and v_j . Then $\{u, v\} \neq \{a_i, a_j\}$ since $a_i a_j \in E$ and $uv \in \overline{E}$. Therefore, we may assign either u or v (or both) as colour to x_{uv} . ◁

Since MULTICOLOURED CLIQUE is W[1]-hard, this proves that LIST COLOURING parameterised by tree-partition-width is W[1]-hard. ◀

We remark that the above proof also shows that LIST COLOURING parameterised by vertex cover is W[1]-hard.

6 Conclusion

In this paper, we combined combinatorial insights and algorithmic tricks to give a space-efficient colouring algorithm.

By combining Logspace Bodlaender’s theorem [14], Lemma 5 and Lemma 4, LIST COLOURING can be solved non-deterministically on graphs of pathwidth k in $O(k \log n)$ space and on graphs of treewidth k in $O(k \log^2 n)$ space.⁴ However, we already do not know the answer to the following question.

▶ **Problem 11.** *Can a non-deterministic Turing machine solve LIST COLOURING for n -vertex graphs of treewidth 2 using $o(\log^2 n)$ space?*

Another natural way to extend trees is by considering graphs of bounded treedepth. Such graphs then also have bounded pathwidth (but the reverse may be false). It has been observed for several problems such as 3-COLOURING and DOMINATING SET that “dynamic programming approaches” (common for pathwidth or treewidth) require space exponential in the width parameter, whereas there is a “branching approach” with space polynomial in

⁴ First compute a tree decomposition $(T, (B_t)_{t \in V(T)})$ of width k for G in $O(\log n)$ space [14], and then compute a (not necessarily optimal) path decomposition of width $O(\log |V(T)|)$ in $O(\log n)$ space for T , and turn this into a path decomposition for G of width $O(k \log n)$ by replacing $t \in V(T)$ with the vertices in its bag B_t . Then use Lemma 4.

treedepth [6]. A simple branching approach also allows LIST COLOURING to be solved in $O(k \log n)$ space on n -vertex graphs of treedepth k . We wonder if the approach in our paper can be adapted to improve this further.

► **Problem 12.** *Can LIST COLOURING be solved in $f(k)g(n) + O(\log n)$ space on graphs of treedepth k , with $g(n) = o(\log n)$ and f a computable function?*

Another interesting direction is what the correct complexity class is for LIST COLOURING parameterised by tree partition width. We do not expect this to be in the W-hierarchy because the required witness size seems to be too large. Moreover, the conjecture [23, Conjecture 2.1] mentioned in the introduction together with Corollary 2 would imply that the problem is not XNLP-hard.

Finally, it would be interesting to study other computational problems than LIST COLOURING. We remark that our results are highly unlikely to generalise to arbitrary Constraint Satisfaction Problems. Recall that there is conjectured to be a $k^* \in \mathbb{N}$ for which LIST COLOURING requires $\omega(\log n)$ space for n -vertex graphs of treewidth k^* . Since LIST COLOURING on n -vertex graphs of treewidth at most k^* can be reduced in logspace to a CSP on at most n variables, each having a lists of size at most n^{k^*} , and binary constraints on the variables, such CSP problems would then also require $\omega(\log n)$ space since k^* is a constant.

References

- 1 Eric Allender, Shiteng Chen, Tiancheng Lou, Periklis A. Papakonstantinou, and Bangsheng Tang. Width-parametrized SAT: Time–space tradeoffs. *Theory of Computing*, 10:297–339, 2014. doi:10.4086/toc.2014.v010a012.
- 2 Sanjeev Arora and Boaz Barak. *Computational Complexity – A Modern Approach*. Cambridge University Press, 2009. URL: <http://www.cambridge.org/catalogue/catalogue.asp?isbn=9780521424264>.
- 3 Hans L. Bodlaender, Carla Groenland, and Hugo Jacob. List colouring trees in logarithmic space. *arXiv*, CoRR(2206.09750), 2022. doi:10.48550/ARXIV.2206.09750.
- 4 Hans L. Bodlaender, Carla Groenland, Jesper Nederlof, and Céline M. F. Swennenhuis. Parameterized problems complete for nondeterministic FPT time and logarithmic space. In *Proceedings 62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021*, pages 193–204, 2021. doi:10.1109/FOCS52979.2021.00027.
- 5 Hans L. Bodlaender, Carla Groenland, and Céline M. F. Swennenhuis. Parameterized complexities of dominating and independent set reconfiguration. In Petr A. Golovach and Meirav Zehavi, editors, *Proceedings 16th International Symposium on Parameterized and Exact Computation, IPEC 2021*, volume 214 of *LIPICs*, pages 9:1–9:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.IPEC.2021.9.
- 6 Li-Hsuan Chen, Felix Reidl, Peter Rossmanith, and Fernando Sánchez Villaamil. Width, depth, and space: Tradeoffs between branching and dynamic programming. *Algorithms*, 11(7):98, 2018. doi:10.3390/a11070098.
- 7 Yijia Chen, Jörg Flum, and Martin Grohe. Bounded nondeterminism and alternation in parameterized complexity theory. In *Proceedings 18th Annual IEEE Conference on Computational Complexity, CCC 2003*, pages 13–29. IEEE Computer Society, 2003. doi:10.1109/CCC.2003.1214407.
- 8 Samir Datta, Nutan Limaye, Prajakta Nimbhorkar, Thomas Thierauf, and Fabian Wagner. Planar graph isomorphism is in log-space. In *Proceedings 24th Annual IEEE Conference on Computational Complexity, CCC 2009*, pages 203–214. IEEE Computer Society, 2009. doi:10.1109/CCC.2009.16.
- 9 Reinhard Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate Texts in Mathematics*. Springer, 2012.
- 10 Guoli Ding and Bogdan Oporowski. Some results on tree decomposition of graphs. *J. Graph Theory*, 20(4):481–499, 1995. doi:10.1002/jgt.3190200412.

- 11 Guoli Ding and Bogdan Oporowski. On tree-partitions of graphs. *Discret. Math.*, 149(1-3):45–58, 1996. doi:10.1016/0012-365X(94)00337-I.
- 12 Rodney G. Downey and Michael R. Fellows. *Parameterized Complexity*. Springer, 1999.
- 13 László Egri, Pavol Hell, Benoît Larose, and Arash Rafiey. Space complexity of list H -colouring: A dichotomy. In Chandra Chekuri, editor, *Proceedings 25th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014*, pages 349–365. SIAM, 2014. doi:10.1137/1.9781611973402.26.
- 14 Michael Elberfeld, Andreas Jakoby, and Till Tantau. Logspace versions of the theorems of Bodlaender and Courcelle. In *Proceedings 51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*, pages 143–152. IEEE Computer Society, 2010. doi:10.1109/FOCS.2010.21.
- 15 Michael Elberfeld, Christoph Stockhusen, and Till Tantau. On the space and circuit complexity of parameterized problems: Classes and completeness. *Algorithmica*, 71(3):661–701, 2015. doi:10.1007/s00453-014-9944-y.
- 16 Michael R. Fellows, Fedor V. Fomin, Daniel Lokshtanov, Frances A. Rosamond, Saket Saurabh, Stefan Szeider, and Carsten Thomassen. On the complexity of some colorful problems parameterized by treewidth. *Inf. Comput.*, 209(2):143–153, 2011. doi:10.1016/j.ic.2010.11.026.
- 17 Naveen Garg, Marina Papatriantafylou, and Philippas Tsigas. Distributed list coloring: How to dynamically allocate frequencies to mobile base stations. In *Proceedings 8th IEEE Symposium on Parallel and Distributed Processing, SPDP 1996*, pages 18–25. IEEE Computer Society, 1996. doi:10.1109/SPDP.1996.570312.
- 18 Sylvain Gravier, Daniel Kobler, and Wieslaw Kubiak. Complexity of list coloring problems with a fixed total number of colors. *Discrete Applied Mathematics*, 117(1-3):65–79, 2002.
- 19 Klaus Jansen and Petra Scheffler. Generalized coloring for tree-like graphs. *Discrete Applied Mathematics*, 75(2):135–155, 1997. doi:10.1016/S0166-218X(96)00085-6.
- 20 Shiva Kintali and Sinziana Munteanu. Computing bounded path decompositions in logspace. *Electron. Colloquium Comput. Complex. (ECCC)*, page 126, 2012. URL: <https://eccc.weizmann.ac.il/report/2012/126>.
- 21 Jan Kynčl and Tomáš Vyskočil. Logspace reduction of directed reachability for bounded genus graphs to the planar case. *ACM Trans. Comput. Theory*, 1(3), 2010. doi:10.1145/1714450.1714451.
- 22 Steven Lindell. A logspace algorithm for tree canonization (extended abstract). In *Proceedings 24th Annual ACM Symposium on Theory of Computing, STOC 1992*, STOC '92, pages 400–404, New York, NY, USA, 1992. Association for Computing Machinery. doi:10.1145/129712.129750.
- 23 Michal Pilipczuk and Marcin Wrochna. On space efficiency of algorithms working on structural decompositions of graphs. *ACM Trans. Comput. Theory*, 9(4):18:1–18:36, 2018. doi:10.1145/3154856.
- 24 Krishna N. Ramachandran, Elizabeth M. Belding-Royer, Kevin C. Almeroth, and Milind M. Buddhikot. Interference-aware channel assignment in multi-radio wireless mesh networks. In *Proceedings 25th IEEE International Conference on Computer Communications, INFOCOM 2006*, pages 1–12. IEEE Computer Society, 2006. doi:10.1109/INFOCOM.2006.177.
- 25 Omer Reingold. Undirected connectivity in log-space. *J. ACM*, 55(4), 2008. doi:10.1145/1391289.1391291.
- 26 Detlef Seese. Tree-partite graphs and the complexity of algorithms. In Lothar Budach, editor, *Proceedings 5th International Conference on Fundamentals of Computation Theory, FCT 1985*, volume 199 of *Lecture Notes in Computer Science*, pages 412–421. Springer, 1985. doi:10.1007/BFb0028825.
- 27 David R. Wood. On tree-partition-width. *Eur. J. Comb.*, 30(5):1245–1253, 2009. doi:10.1016/j.ejc.2008.11.010.