



**Contents**

<b>1</b>	<b>Introduction</b>	<b>4</b>
<b>2</b>	<b>Maintaining Monotone Dynamic Programming Tables</b>	<b>8</b>
<b>3</b>	<b>Fully Dynamic Knapsack</b>	<b>11</b>
3.1	Knapsack via Convolution of Monotone Functions . . . . .	11
3.2	Dynamically Maintaining a Small Instance . . . . .	14
<b>4</b>	<b>Technical Overview</b>	<b>17</b>
4.1	Monotonizing the DP of Feldmann and Foschini . . . . .	17
4.1.1	Computing the DP . . . . .	19
4.2	Inverting the DP of Andreev et al. . . . .	22
<b>A</b>	<b>Organization of the Appendix</b>	<b>26</b>
<b>B</b>	<b>Further Related Work</b>	<b>26</b>
<b>C</b>	<b>Preliminaries</b>	<b>28</b>
C.1	Räcke Tree . . . . .	28
C.2	Okay-Behaved DPs . . . . .	29
<b>D</b>	<b>Balanced Graph Partitioning</b>	<b>30</b>
D.1	The Exact DP . . . . .	31
D.1.1	DP Definition . . . . .	33
D.1.2	Computing the DP . . . . .	33
D.2	The Approximate DP . . . . .	36
D.3	Computing the Result . . . . .	40
D.4	Extension to General Graphs . . . . .	43
D.5	Extension to the Dynamic Setting . . . . .	45
<b>E</b>	<b>Simultaneous Source Location</b>	<b>48</b>
E.1	The Exact DP . . . . .	49
E.1.1	DP Definition . . . . .	49
E.1.2	Computing the DP . . . . .	51
E.2	The Approximate DP . . . . .	53
E.3	Extension to General Graphs (Proof of Theorem 4) . . . . .	54
E.4	Extension to the Dynamic Setting (Proof of Theorem 5) . . . . .	54
<b>F</b>	<b>Recourse Bounds</b>	<b>56</b>
F.1	Proof of Theorem 34 . . . . .	57
F.2	Proof of Theorem 35 . . . . .	59
<b>G</b>	<b>Non-Monotone Functions and <math>\ell_\infty</math>-Necklace Alignment</b>	<b>60</b>
G.1	Piecewise Constant Functions With Non-Monotonicities . . . . .	60
G.1.1	Proof of Lemma 40 . . . . .	62
G.2	$\ell_\infty$ -Necklace Alignment . . . . .	64
G.2.1	The Static Algorithm . . . . .	65
G.2.2	The Dynamic Algorithm . . . . .	68

<b>H Omitted Proofs</b>	<b>69</b>
H.1 Proof of Lemma 6 . . . . .	69
H.2 Proof of Lemma 7 . . . . .	70
H.3 Proof of Theorem 9 . . . . .	71
H.4 Proof of Theorem 10 . . . . .	72
H.5 Property of the Räcke Tree . . . . .	72
H.6 Proof of Lemma 18 . . . . .	73
H.7 Proof of Lemma 19 . . . . .	73
H.8 Proof of Lemma 20 . . . . .	73
H.9 Proof of Lemma 21 . . . . .	73
H.10 Proof of Lemma 22 . . . . .	74



## 1 Introduction

Dynamic programming (DP) is one of the fundamental paradigms in algorithm design. In the DP paradigm, a complex problem is broken up into simpler subproblems and then the original problem is solved by combining the solutions for the subproblems. One of the drawbacks of DP algorithms is that in practice they are often slow and memory-intensive: for inputs of size  $n$  their running time is typically  $\Omega(n^2)$ , and when the DP table is stored using a two-dimensional array they also need space  $\Omega(n^2)$ .

This motivated researchers to develop more efficient DP algorithms with near-linear time and space. Indeed, such improvements are possible under a wide range of conditions on the DP tables [2, 12, 19, 22, 31, 35, 45, 48, 49, 64], such as the Monge property, total monotonicity, certain convexity and concavity properties, or the Knuth–Yao quadrangle-inequality; we discuss these properties in more detail in Appendix B. When these properties hold, typically one does not have to compute the entire DP table but instead only has to compute  $O(n)$  DP entries which reveal the optimal solution.

However, we are not aware of any property for DPs that yields efficient *dynamic* algorithms, i.e., algorithms that provide efficient update operations when the input changes. One might find this somewhat surprising because, from a conceptual point of view, many dynamic algorithms hierarchically partition the input and maintain solutions for subproblems; this is quite similar to how many DP schemes are derived. Indeed, this conceptual similarity is exploited by many “hand-crafted” algorithms (e.g., [26, 38]) which start with a DP scheme and then show how to maintain it dynamically under input changes. However, such algorithms are often quite involved and their analysis often requires sophisticated charging schemes.

Hence, it is natural to ask whether there exists a *general criterion* which, if satisfied, guarantees that a given DP can be updated efficiently under input changes.

**Our Contributions.** The main contribution of our paper is the introduction of a general criterion which allows to approximate all entries of a DP table up to a factor of  $1 + \epsilon$ . We show that if our criterion is satisfied by a DP (with suitable parameters) then:

- In the dynamic setting, we can maintain a  $(1 + \epsilon)$ -approximation of the entire DP table using polylogarithmic update time (see Theorem 10).
- In the static setting, we can compute a  $(1 + \epsilon)$ -approximation of the DP table in near-linear time and space (see Theorem 9).

Our criterion essentially asserts that the *rows* of the DP tables should be *monotone* and that the dependency graph of the DP should be a DAG, where the sets of reachable nodes are small, among some other technical conditions (see Definition 8 for the formal definition). Our criterion is incomparable to the Monge property, total monotonicity or other criteria from the literature (see Appendix B for a more detailed discussion).

To obtain our results, we introduce a novel data structure for maintaining DPs which satisfy our criterion. Our data structure is based on the idea of storing the DP rows using *monotone piecewise constant* functions. The monotonicity of the DP rows will allow us to ensure that our functions only contain very few pieces. Then we show that we can perform operations on such functions very efficiently, with the running times only depending on the number of pieces. This is crucial because it allows us to compute an *entire*  $(1 + \delta)$ -approximate DP row in time just  $\text{polylog}(W)$ , even when the DP has  $\Omega(n)$  columns, assuming that the DP entries are from  $[0, W]$ . Note that if  $W \leq \text{poly}(n)$  then this decreases the running time for computing an entire row from  $\Omega(n)$  to just  $\text{polylog}(n)$ . Additionally, this also allows us to store each row *using only*  $\text{polylog}(W)$  *space* rather than storing it in an array of size  $\Omega(n)$ . We present our criterion and the details of our data structure in Section 2.

As applications of our data structure, we obtain new static and dynamic algorithms for various problems. We present new algorithms for *k-balanced partitioning*, *simultaneous source location* and for *fully dynamic knapsack*. Next, we describe these results in detail; we discuss more related work in Appendix B.

**Our Results for Fully Dynamic 0-1 Knapsack.** First, we provide a novel algorithm for fully dynamic 0-1 knapsack. In this problem, the input consists of a knapsack size  $B \in \mathbb{R}_+$  and a set of  $n$  items, where each item  $i \in [n]$  has a weight  $w_i \in \mathbb{R}_+$  and a price  $p_i \in [1, \infty)$ . The goal is to find a set of items  $I$  that maximizes  $\sum_{i \in I} p_i$  while satisfying the constraint  $\sum_{i \in I} w_i \leq B$ . In the dynamic version of the problem, items are inserted and deleted. More concretely, we consider the following update operations: *insert*( $p_i, w_i$ ), in which the price and weight of item  $i$  are set to  $p_i \in [1, \infty)$  and  $w_i \in \mathbb{R}_+$ , respectively, and *delete*( $i$ ), where item  $i$  is removed from the set of items.

Our main result is a dynamic  $(1 + \epsilon)$ -approximation algorithm with worst-case update time  $\epsilon^{-2} \cdot \log^2(nW) \cdot \text{polylog}(1/\epsilon, \log(nW))$ , where  $W = \sum_i p_i$ . Our algorithm improves upon a recent result by Eberle, Megow, Nölke, Simon and Wiese [29] that also maintained a  $(1 + \epsilon)$ -approximate solution with update time  $O(\epsilon^{-9} \log^4(nW))$ .

► **Theorem 1.** *Let  $\epsilon > 0$ . There exists an algorithm for fully dynamic knapsack that maintains a  $(1 + \epsilon)$ -approximate solution with worst-case update time  $\frac{1}{\epsilon^2} \log^2(nW) \text{polylog}(\frac{1}{\epsilon} \log(nW))$ .*

We will also show that we can return the maintained solution  $I$  in time  $O(|I|)$  and that we can answer queries whether a given item  $i \in [n]$  is contained in  $I$  in time  $O(1)$ . This matches the query times of [29].

To obtain this result, we first derive a slightly slower algorithm as a simple application of our data structure for maintaining DPs with monotone rows. Then we use this algorithm together with additional ideas to obtain the theorem (see Section 3).

Since our dynamic algorithm is based on a DP, it is possible that the solution changes significantly after each update. However, in the appendix (Theorem 34) we prove a lower bound, showing that every dynamic  $(1 + \epsilon)$ -approximation algorithm for knapsack must either make a lot of changes to the solution after each update or store many (potentially substantially different) solutions between which it can switch after each update. This implies that maintaining a single explicit solution with polylogarithmic update times is not possible and the property of our algorithm cannot be avoided.

**Our Results for *k*-Balanced Partitioning.** Our most technically challenging result is for *k-balanced graph partitioning*. In this problem, the input consists of an integer  $k$  and an undirected weighted graph  $G = (V, E, \text{cap})$  with  $n$  vertices, where  $\text{cap} : E \rightarrow W_\infty$  is a weight function on the edges with weights in  $W_\infty := [1, W] \cup \{0, \infty\}$ . The goal is to find a partition  $V_1, \dots, V_k$  of the vertices such that  $|V_i| \leq \lceil |V|/k \rceil$  for all  $i$  and the weight of the edges which are cut by the partition is minimized. More formally, we want to minimize  $\text{cut}(V_1, \dots, V_k) := \sum_{i=1}^k \sum_{\{u,v\} \in E \cap (V_i \times (V \setminus V_i))} \text{cap}(u, v)$ .

We note that this problem is highly relevant in theory [5, 32–34] and in practice [18, 28, 44, 55], where algorithms for balanced graph partitioning are often used as a preprocessing step for large scale data analytics. Obtaining practical improvements for this problem is of considerable interest in applied communities [18] and, for instance, the popular METIS heuristic [44] has 1,400+ citations.

Since the above problem is NP-hard to approximate within a factor of  $n^{1-\epsilon}$  for any  $\epsilon > 0$  even on trees [34], we consider bicriteria approximation algorithms. Given an undirected weighted graph  $G = (V, E, \text{cap})$ , a partition  $V_1, \dots, V_k$  of  $V$  is a *bicriteria*  $(\alpha, \beta)$ -approximate solution if  $|V_i| \leq \beta \lceil n/k \rceil$  for all  $i$  and  $\text{cut}(V_1, \dots, V_k) \leq \alpha \cdot \text{cut}(\text{OPT})$ , where  $\text{OPT} =$

$(V_1^*, \dots, V_k^*)$  is the optimal solution with  $|V_i^*| \leq \lceil n/k \rceil$  for all  $i$ . We note that the previously mentioned hardness result implies that any algorithm that computes a bicriteria  $(\alpha, 1 + \epsilon)$ -approximation for any  $\alpha \geq 1$  and whose running time depends only polynomially on  $n$ , must have a running time depending super-polynomially on  $1/\epsilon$ , unless  $\mathbf{P} = \mathbf{NP}$ .<sup>1</sup>

Our main result for the static setting is presented in the following theorem. It gives the first algorithm with polylogarithmic approximation ratio for this problem *with near-linear running time*. More concretely, we compute a bicriteria  $(O(\log^4 n), 1 + \epsilon)$ -approximation in near-linear time for constant  $k$ . For comparison, the best approximation ratio achieved by a polynomial-time algorithm [34] is a bicriteria  $(O(\log^{1.5} n \log \log n), 1 + \epsilon)$ -approximation with running time  $\Omega(n^4)$ .

► **Theorem 2.** *Let  $\epsilon > 0$  and  $k \in \mathbb{N}$ . Let  $G = (V, E, \text{cap})$  be an undirected weighted graph with  $n$  vertices and  $m$  edges and edge weights in  $W_\infty$ . Then for the  $k$ -balanced partition problem we can compute:*

- An  $(O(\log^4 n), 1 + \epsilon)$ -approximation in time  $(k/\epsilon)^{O(\log(1/\epsilon)/\epsilon)} \cdot O'(m \cdot \log^2(W)) + (k/\epsilon)^{O(1/\epsilon^2)}$ .<sup>2</sup>
- A  $(1 + \epsilon, 1 + \epsilon)$ -approximation in time  $(k/\epsilon)^{O(\log(1/\epsilon)/\epsilon)} \cdot O'(n \cdot h^2 \cdot \log^2(W)) + (k/\epsilon)^{O(1/\epsilon^2)}$  if  $G$  is a tree of height  $h$ .
- A  $(1, 1 + \epsilon)$ -approximation in time  $(k/\epsilon)^{O(\log(1/\epsilon)/\epsilon)} \cdot O'(n^4 \cdot \log^2(W)) + (k/\epsilon)^{O(1/\epsilon^2)}$  if  $G$  is a tree.

Furthermore, we extend our results to the dynamic setting in which the graph  $G$  is undergoing edge insertions and deletions. In the following theorem, we present the *first dynamic algorithm with subpolynomial update time* for this problem. We again consider bicriteria approximation algorithms with update and query times depending super-polynomially on  $1/\epsilon$ ; this cannot be avoided since if we computed  $(\alpha, 1)$ -approximations for any  $\alpha \geq 1$  or if we had a polynomial dependency on  $1/\epsilon$ , then the hardness result from above implies that our update and query times must be super-polynomial in  $n$  (unless  $\mathbf{P} = \mathbf{NP}$ ).

► **Theorem 3.** *Let  $\epsilon > 0$  and  $k \in \mathbb{N}$ . Let  $G = (V, E, \text{cap})$  be an undirected weighted graph with  $n$  vertices that is undergoing edge insertions and deletions. Then for the  $k$ -balanced partition problem we can maintain:*

- An  $(n^{o(1)}, 1 + \epsilon)$ -approximate solution with amortized update time  $(k/\epsilon)^{O(\log(1/\epsilon)/\epsilon)} \cdot n^{o(1)} \cdot O'(\log^2(W))$  and query time  $(k/\epsilon)^{O(1/\epsilon^2)}$  if  $G$  is unweighted.
- A  $(1 + \epsilon, 1 + \epsilon)$ -approximate solution with worst-case update time  $(k/\epsilon)^{O(\log(1/\epsilon)/\epsilon)} \cdot O'(h^3 \cdot \log^2(W))$  and query time  $(k/\epsilon)^{O(1/\epsilon^2)}$  if  $G$  is a tree of height  $h$ .

Our approach is inspired by the DP of Feldmann and Foschini [34]. However, the DP rows in the algorithm of [34] are not monotone and, hence, their DP cannot directly be sped up by our approach. Therefore, we first simplify and generalize the exact DP of Feldmann and Foschini to make it monotone. The DP we obtain eventually is still slightly too complex to fit into our black-box framework, but we show that the ideas from our framework can still be used to obtain the result. In Section 4.1, we provide a technical overview.

Again, it is possible that the solution maintained by our algorithm changes substantially after each update. Similar to above we show in the appendix (Theorem 35) that this cannot be avoided when considering subpolynomial update times.

<sup>1</sup> If we had an algorithm that computes a bicriteria  $(\alpha, 1 + \epsilon)$ -approximation in time  $\text{poly}(n, 1/\epsilon)$  then we could set  $\epsilon = 1/(2n)$  which implies that all partitions have size  $\lceil n/k \rceil$ . Thus we can compute a bicriteria  $(\alpha, 1)$ -approximate solution in time  $\text{poly}(n)$  which contradicts the hardness result, unless  $\mathbf{P} = \mathbf{NP}$ .

<sup>2</sup> We use the notation  $O'(\cdot)$  to suppress factors in  $\text{poly}(\log n, k, \log(1/\epsilon), \log \log(W))$ .

**Our Results for Simultaneous Source Location.** Next, we provide efficient algorithms for the simultaneous source location problem by Andreev, Garrod, Golovin, Maggs and Meyerson [4]. In this problem, the input consists of an undirected graph  $G = (V, E, \text{cap}, d)$  with a capacity function  $\text{cap}: E \rightarrow W_\infty$  on the edges and a demand function  $d: V \rightarrow W_\infty$  on the vertices. The goal is to select a minimum set  $S \subseteq V$  of *sources* that can simultaneously supply all vertex demands. More concretely, a set of sources  $S$  is *feasible* if there exists a flow from the vertices in  $S$  that supplies demand  $d(v)$  to all vertices  $v \in V$  and that does not violate the capacity constraints on the edges. The objective is to find a feasible set of sources of minimum size.

We will again consider bicriteria approximation algorithms. Let  $S^*$  be the optimal solution for the simultaneous source location problem. Then we say that  $S$  is a *bicriteria*  $(\alpha, \beta)$ -*approximate solution* if  $|S| \leq \alpha |S^*|$  and if  $S$  is a feasible set of sources when all edge capacities are increased by a factor  $\beta$ .

The following theorem summarizes our main results. It presents *the first near-linear time algorithm* for simultaneous source location that computes a  $(1 + \epsilon)$ -approximate solution while only exceeding the edge capacities by a  $O(\log^4 n)$  factor. In comparison, the best algorithm with arbitrary polynomial processing time computes a bicriteria  $(1, O(\log^2 n \log \log n))$ -approximate solution in time  $\Omega(n^3)$  [4].

► **Theorem 4.** *Let  $\epsilon > 0$ . Let  $G = (V, E, \text{cap}, d)$  be an undirected weighted graph with  $n$  vertices and  $m$  edges. Then for the simultaneous source location problem we can compute:*

- *A  $(1 + \epsilon, O(\log^4(n)))$ -approximation in time<sup>3</sup>  $\tilde{O}(\frac{1}{\epsilon^2}m)$ .*
- *A  $(1 + \epsilon, 1)$ -approximation in time  $\tilde{O}(\frac{1}{\epsilon^2}h^2 \cdot n)$  if  $G$  is a tree of height  $h$ .*

Next, we turn to dynamic versions of the problem. We consider the following update operations: *SetDemand* $(v, d)$ : updates the demand of vertex  $v$  to  $d(v) = d$ , *SetCapacity* $((u, v), c)$ : updates the capacity of the edge  $(u, v)$  to  $\text{cap}(u, v) = c$ , *Remove* $(u, v)$ : removes the edge  $(u, v)$ , *Insert* $((u, v), c)$ : inserts the edge  $(u, v)$  with capacity  $\text{cap}(u, v) = c$ .

We obtain *the first dynamic algorithms with subpolynomial update times* for this problem, which exceed the edge capacities only by a small subpolynomial factor.

► **Theorem 5.** *Let  $\epsilon > 0$ . Let  $G = (V, E, \text{cap}, d)$  be a graph with  $n$  vertices and  $m$  edges that is undergoing the update operations given above. Then for the simultaneous source location problem we can maintain:*

- *A  $(1 + \epsilon, n^{o(1)})$ -approximation with amortized update time  $n^{o(1)}/\epsilon^2$  and preprocessing time  $O(n^2/\epsilon^2)$  if all edge capacities are 1.*
- *A  $(1 + \epsilon, O(\log^4(n)))$ -approximation with worst-case update time  $\tilde{O}(1/\epsilon^2)$  and preprocessing time  $\tilde{O}(m)$  if we only allow the update operation *SetDemand* $(v, d)$ .*
- *A  $(1 + \epsilon, O(\log^2(n) \log \log(n)))$ -approximation with worst-case update time  $\tilde{O}(1/\epsilon^2)$  and preprocessing time  $\text{poly}(n)$  if we only allow the update operation *SetDemand* $(v, d)$ .*
- *A  $(1 + \epsilon, 1)$ -approximate solution with worst-case update time  $\tilde{O}(h^3/\epsilon^2)$  and preprocessing time  $O(n^2/\epsilon^2)$  if  $G$  is a tree of height  $h$ .*

To obtain these results, we use a similar DP approach as the one used by Andreev et al. [4]. Interestingly, the DP function that we use essentially computes *the inverse function* of the one used by Andreev et al. We sketch the details of this approach in Section 4.2. After making these changes, the theorems become straightforward applications of our data structure for maintaining DPs with monotone rows.

<sup>3</sup> We write  $\tilde{O}(f(n, \epsilon, W))$  to denote running times of the form  $f(n, \epsilon, W) \cdot \text{polylog}(n, \epsilon, \log W)$ .

**Organization of Our Paper.** In Section 2 we provide the details of our condition for DPs with monotone rows. In Section 3 we present our results for 0-1 Knapsack which nicely illustrate the applicability of our black-box framework from Section 2. We provide a technical overview of our more involved results for  $k$ -Balanced Graph Partitioning and for Simultaneous Source Location in Section 4. We give an overview of the appendix in Appendix A. In the appendix we also present more related work and the full proofs of our results. We present omitted proofs from the main text in Appendix H.

**Open Problems and Future Work.** In the future, it will be interesting to use our framework to obtain more dynamic algorithms based on existing DPs. We believe that this is interesting both in theory and in practice. Furthermore, it is intriguing to ask whether our criterion from Definition 8 can be generalized. Indeed, our approach was built around approximating monotone functions using piecewise constant functions, which can be viewed as piecewise degree-0 polynomials. An interesting question is whether we can obtain a more general criterion if we approximate DP rows using pieces of higher-degree polynomials, such as splines. Results in this direction might be possible; for example, in Appendix G we give a side result for the case when the functions contain a small number of non-monotonicities and derive a dynamic algorithm for the  $\ell_\infty$ -necklace problem.

## 2 Maintaining Monotone Dynamic Programming Tables

In this section, we introduce our notion of DP tables with *monotone rows* and the additional technical assumptions that we are making. Then we present our data structure for efficiently maintaining DP tables that satisfy our assumptions. In our data structure, we will store the rows of the DP using piecewise constant functions, which we will introduce first.

**List Representation of Piecewise Constant Functions.** Let  $t \in \mathbb{R}$ ,  $W \in [1, \infty)$  and set  $W_\infty := \{0\} \cup [1, W] \cup \{+\infty\}$ . A function  $f: [0, t] \rightarrow W_\infty$  is *piecewise constant with  $p$  pieces* if there exist real numbers  $0 = x_0 < x_1 < x_2 < \dots < x_p = t$  and numbers  $y_1, \dots, y_p \in W_\infty$  such that on each interval  $[x_{i-1}, x_i)$ ,  $f$  is constant and has value  $y_i$ . More formally, for all  $i \in \{1, \dots, p\}$  we have  $f(x) = y_i$  for all real numbers  $x \in [x_{i-1}, x_i)$  and  $f(x_p) = y_p$ . Note that we need the condition  $f(x_p) = y_p$  such that  $f$  is defined on the whole domain.

In the *list representation* of a piecewise constant function  $f$ , we use a doubly linked list to store the pairs  $(x_1, y_1), \dots, (x_p, y_p)$ . We also store the pairs  $(x_i, y_i)$  in a binary search tree that is sorted by the  $x_i$ -values, which allows us to compute a function value  $f(x)$  in time  $O(\log p)$  for all  $x \in [0, t]$ . In the following, we assume that all piecewise constant functions we consider are stored in the list representation with an additional binary search tree.

One of the main observations we use is that many operations on piecewise constant functions are efficient if there are only few pieces. The following lemma shows that several operations can be computed in time almost linear in the number of pieces of the function, rather than in time depending on the size of the domain of  $f$ .<sup>4</sup> For  $\delta > 0$  and  $y \in W_\infty$ , we write  $\lceil y \rceil_{1+\delta}$  to denote the smallest power of  $1+\delta$  that is at least  $y$ , i.e.,  $\lceil y \rceil_{1+\delta} = \min\{(1+\delta)^i : (1+\delta)^i \geq y, i \in \mathbb{N}\}$ ; we follow the convention that  $\lceil 0 \rceil_{1+\delta} = 0$  and  $\lceil \infty \rceil_{1+\delta} = \infty$ .

► **Lemma 6.** *Let  $t \in \mathbb{R}$  and  $c \in \mathbb{R}_+$ . Let  $g, h: [0, t] \rightarrow W_\infty$  be monotone and piecewise constant functions with  $p_g$  and  $p_h$  pieces, resp. Then we can compute the following functions:*

- $f_{\min}(x) := \min\{g(x), h(x)\}$  with at most  $p_g + p_h$  pieces in time  $O((p_g + p_h) \log(p_g + p_h))$ ;

<sup>4</sup> We note that computing the operations themselves can be done in linear time. However, since we also store the pairs  $(x_i, y_i)$  of the list representations in a binary search tree, the running times in the lemma include an additional logarithmic factor.



- $f_{\text{shift}}(x) := g(x - c)$  for  $x \geq c$ ,  $f_{\text{shift}}(x) = g(0)$  for  $x < c$  with at most  $p_g$  pieces in time  $O(p_g \log(p_g))$ ;
- $f_{\text{add}}(x) := g(x) + h(x)$ , with at most  $p_g + p_h$  pieces in time  $O((p_g + p_h) \log(p_g + p_h))$ ;
- $f_{\text{round}}(x) := \lceil g(x) \rceil_{1+\delta}$  for  $\delta > 0$  with at most  $2 + \lceil \log_{1+\delta}(W) \rceil$  pieces in time  $O(p_g \log(p_g))$ .

Note that if we set  $\tilde{f} = \lceil f \rceil_{1+\delta}$  then  $\tilde{f}$  is a  $(1 + \delta)$ -approximation of  $f$  in the following sense. For  $\alpha > 1$ , we say that a function  $\tilde{f}: [0, t] \rightarrow W_\infty$   $\alpha$ -approximates a function  $f: [0, t] \rightarrow W_\infty$  if for all  $x \in [0, t]$ ,

$$f(x) \leq \tilde{f}(x) \leq \alpha \cdot f(x). \quad (1)$$

Furthermore, if  $f$  is monotone then the rounded function  $\tilde{f}$  contains at most  $O(\log_{1+\delta}(W))$  pieces. This will be crucial later because this ensures that, if we perform a single rounding operation for each row of our DP table, the resulting function will have few pieces and operations on the function can be performed efficiently.

Next, consider functions  $f_1, f_2: [0, t] \rightarrow W_\infty$ . A function  $f: [0, t] \rightarrow W_\infty$  is the  $(\min, +)$ -convolution  $f_1 \oplus f_2$  if for all  $x \in [0, t]$ ,  $f(x) = (f_1 \oplus f_2)(x) := \min_{\bar{x} \in [0, x]} f_1(\bar{x}) + f_2(x - \bar{x})$ . Such convolutions are highly useful for the computation of many DPs. The following lemma shows that we can efficiently compute the convolution of piecewise constant functions.

► **Lemma 7.** *Let  $f_1, f_2: [0, t] \rightarrow W_\infty$  be piecewise constant functions with at most  $p$  pieces and assume that one of them is monotonically decreasing. Then we can compute the function  $f: [0, t] \rightarrow W_\infty$  with  $f = f_1 \oplus f_2$  in time  $O(p^2 \log p)$  and  $f$  is a piecewise constant function with  $O(p^2)$  pieces. Furthermore, after computing  $f$ , for any  $x \in [0, t]$  we can return a value  $\bar{x}^* \in [0, t]$  such that  $f(x) = f_1(\bar{x}^*) + f_2(x - \bar{x}^*)$  in time  $O(\log p)$ .*

Now observe that Lemma 7 has a drawback for our approach: The number of pieces (i.e., the complexity of the functions) grows quadratically with every application. An important property which can be used to mitigate this issue is that the result of the convolution is still a monotone function, as we show in Lemma 22 in the appendix. Later, to keep the number of pieces in our functions small, after each convolution that we perform via Lemma 7 (and that might grow the number of pieces quadratically), we perform a rounding operation  $\lceil \cdot \rceil_{1+\delta}$  (see Lemma 6). This loses a factor  $1 + \delta$  in approximation but guarantees that the resulting function has  $O(\log_{1+\delta}(W))$  pieces. This will be crucial to ensure that our functions have only few pieces.

**Maintaining DPs With Monotone Rows.** Next, we introduce our DP scheme formally. We consider DP tables with a finite set of rows  $\mathcal{I}$  and a set of columns  $\mathcal{J}$ , with entries taking values in  $W_\infty$ . We will consider DP tables as functions  $\text{DP}: \mathcal{I} \times \mathcal{J} \rightarrow W_\infty$ .<sup>5</sup> Further, we will associate the  $i$ 'th row of the DP with a function  $\text{DP}(i, \cdot): \mathcal{J} \rightarrow W_\infty$ , and we store each such function  $\text{DP}(i, \cdot)$  using piecewise constant functions from above.

Next, we introduce the dependency graph for the rows of our DP. More concretely, the dependency graph  $D = (\mathcal{I}, E_D)$  is a directed graph that has the rows  $\mathcal{I}$  as vertices and a directed edge  $(i', i)$  between two rows if for some columns  $j, j' \in \mathcal{J}$  the entry  $\text{DP}(i', j')$  is required to compute  $\text{DP}(i, j)$ . We write  $\text{In}(i) = \{i' \in \mathcal{I}: (i', i) \in E_D\}$  to denote the set of rows  $i'$  that are required to compute row  $i$ . For the rest of the paper we will assume that the dependency graph is a DAG, which is the case for all applications that we study. We will also write  $\text{Reach}(i)$  to denote the set of vertices that are reachable from row  $i$  in  $D$ .

<sup>5</sup> Even though our definition may suggest that we only consider two-dimensional DP tables, we do not require an order on  $\mathcal{I}$  and we allow  $\mathcal{I}$  to be any finite set. For example, in Section D we will set  $\mathcal{I}$  to 3-tuples corresponding to the parameters of a four-dimensional DP.

Since we assume that the dependency graph is a DAG, we can compute the  $i$ 'th DP row as soon as we have computed the solutions for the DP rows in  $\text{In}(i)$ . We assume that this is done via a *procedure*  $\mathcal{P}_i$  that takes as input the DP rows  $\text{DP}(i', \cdot)$  for all  $i' \in \text{In}(i)$  and returns the row  $\text{DP}(i, \cdot) = \mathcal{P}_i(\{\text{DP}(i', \cdot) : i' \in \text{In}(i)\})$ .

Next, we come to our condition which encodes when our scheme applies. In the definition and for the rest of the paper, we write ADP to refer to an approximate DP table, which approximates the exact DP table DP. Let  $\beta > 1$ . We say that ADP  $\beta$ -approximates DP if  $\text{DP}(i, j) \leq \text{ADP}(i, j) \leq \beta \text{DP}(i, j)$  for all  $i \in \mathcal{I}, j \in \mathcal{J}$ .

► **Definition 8.** *A DP table is  $(h, \alpha, p)$ -well-behaved if it satisfies the following conditions:*

1. *(Monotonicity:) For all  $i \in \mathcal{I}$ , the function  $\text{DP}(i, \cdot)$  is monotone.*
2. *(Dependency graph:) The dependency graph is a DAG and  $|\text{Reach}(i)| \leq h$  for all  $i \in \mathcal{I}$ .*
3. *(Sensitivity:) Suppose  $\beta > 1$  and for all  $i' \in \text{In}(i)$ , we obtain a  $\beta$ -approximation  $\text{ADP}(i', \cdot)$  of  $\text{DP}(i', \cdot)$ . Then applying  $\mathcal{P}_i$  on the  $\text{ADP}(i', \cdot)$  yields a  $\beta$ -approximation of  $\text{DP}(i, \cdot)$ , i.e.,*

$$\text{DP}(i, \cdot) \leq \mathcal{P}_i(\{\text{ADP}(i', \cdot) : i' \in \text{In}(i)\}) \leq \beta \cdot \text{DP}(i, \cdot).$$

4. *(Pieces:) For each procedure  $\mathcal{P}_i$  there exists an approximate procedure  $\tilde{\mathcal{P}}_i$  such that:*
  - (a)  $\tilde{\mathcal{P}}_i(\{\text{ADP}(i', \cdot) : i' \in \text{In}(i)\})$  *is an  $\alpha$ -approximation of  $\mathcal{P}_i(\{\text{ADP}(i', \cdot) : i' \in \text{In}(i)\})$ ,*
  - (b)  $\tilde{\mathcal{P}}_i$  *can be computed as the composition of a constant number of operations from Lemma 6 and at most one application of Lemma 7, and*
  - (c)  $\tilde{\mathcal{P}}_i$  *returns a monotone piecewise constant function with at most  $p$  pieces.*

The definition is motivated in the following way: our operations on the piecewise constant functions have efficient running times when the functions are monotone and have few pieces. This is ensured by Properties (1), 4(b), and 4(c). Next, rounding errors cannot compound too much if each row can only reach  $h$  other rows and the sensitivity condition is satisfied. This is ensured by Properties (2), (3), and 4(a).

Even though the definition might look slightly technical at first glance, it applies in many settings. In particular, Property (2) is satisfied when the dependency graph is a rooted tree of height  $h$  in which all edges point towards the root; this is the case in all of our applications. The other conditions are immediately satisfied by our DP for 0-1 Knapsack in Section 3 and the DP for simultaneous source location in Section E. However, our DP for balanced graph partitioning violates Property (4b) of Definition 8. Hence, we will also consider a weaker assumption in Section C.2 which, however, will not allow for nice black-box results, such as Theorems 9 and 10 below.

Next, we state our main results. They imply that we obtain static  $(1 + \epsilon)$ -approximation algorithms running in near-linear time and space for  $(\tilde{O}(1), \ln(1 + \epsilon)/\tilde{O}(1), \tilde{O}(1))$ -well-behaved DPs. They also show that under this assumption, we can dynamically maintain  $(1 + \epsilon)$ -approximate DP solutions with polylogarithmic update times.

Our main theorem for static algorithms is as follows.

► **Theorem 9.** *Consider an  $(h, \alpha, p)$ -well-behaved DP. Then we can compute an approximate DP table ADP which  $\alpha^{h+1}$ -approximates DP in time and space  $O(|\mathcal{I}| \cdot p^2 \log(p))$ .*

Later, we will apply the theorem to DPs with dependency trees of logarithmic heights  $h = O(\log n)$ , we will set the approximation ratio to  $\alpha = \ln(1 + \epsilon)/(h + 1)$ , and the number of pieces to  $p = \text{polylog}(W)$ . This will yield our desired algorithms with near-linear running time  $\tilde{O}(|\mathcal{I}|)$  and space usage. Note that this is a big improvement upon the brute-force running times and space usages of  $\Omega(|\mathcal{I}| \cdot |\mathcal{J}|)$ .

The proof of the theorem follows from observing that when moving from one vertex to another in the dependency graph, we lose a multiplicative  $\alpha$ -factor in the approximation ratio; as each vertex can only reach  $h$  other vertices, this will compound to at most  $\alpha^{h+1}$ . Combining the assumptions on the functions  $\tilde{\mathcal{P}}_i$  and the results from Lemmas 6 and 7, we get that each row  $\text{ADP}(i, \cdot)$  can be computed in time  $O(p^2 \log(p))$  which gives  $O(|I| \cdot p^2 \log(p))$  total time.

We also give the following extension to the dynamic setting which shows that if one of the DP rows changes, we can update *the entire table* efficiently.

► **Theorem 10.** *Consider an  $(h, \alpha, p)$ -well-behaved DP and suppose that row  $i$  is changed. Then we can update our approximate DP table  $\text{ADP}$  such that after time  $O(h \cdot p^2 \log(p))$  it is an  $\alpha^{h+1}$ -approximation of DP.*

As before, we will typically use the theorem with  $h = O(\log n)$ ,  $\alpha = \ln(1 + \epsilon)/(h + 1)$  and  $p = \text{polylog}(W)$ . This will then result in our desired polylogarithmic update times. Note that this is a significant speedup compared to storing the DP tables using two-dimensional arrays: in that case even updating *a single row* would take time  $\Omega(|\mathcal{J}|)$ , which in many applications would already be linear in the size of the input.

The theorem follows from observing that after a row  $i$  changes, we only have to update those rows which can be reached from  $i$  in the dependency graph. But these can be at most  $h$  and each of them can be updated in time  $O(p^2 \log(p))$  by Lemmas 6 and 7.

### 3 Fully Dynamic Knapsack

In *0-1 knapsack*, the input consists of a knapsack size  $B \in \mathbb{R}_+$  and a set of  $n$  items, where each item  $i \in [n]$  has a weight  $w_i \in \mathbb{R}_+$  and a price  $p_i \in [1, \infty)$ . The goal is to find a set of items  $I$  that maximizes  $\sum_{i \in I} p_i$  while satisfying the constraint  $\sum_{i \in I} w_i \leq B$ . For a set of items  $I \subseteq [n]$ , we refer to the sum  $\sum_{i \in I} w_i$  as *the weight of  $I$* .

For the rest of this section we set  $W = \sum_i p_i$  and  $t = \sum_{i \in [n]} w_i$ .

Next, we first derive a dynamic algorithm with update time  $\tilde{O}(\log^3(n) \log^2(W)/\epsilon^2)$  which is based on our framework for DPs with monotone rows. Then we will use this algorithm as a subroutine to obtain a faster algorithm with update time  $\tilde{O}(\log^2(nW)/\epsilon^2)$  in Section 3.2; this will prove Theorem 1.

► **Theorem 1.** *Let  $\epsilon > 0$ . There exists an algorithm for fully dynamic knapsack that maintains a  $(1 + \epsilon)$ -approximate solution with worst-case update time  $\frac{1}{2} \log^2(nW) \text{polylog}(\frac{1}{\epsilon} \log(nW))$ .*

Below we will also show that we can return the maintained solution  $I$  in time  $O(|I|)$  and that we answer queries whether a given item  $i \in [n]$  is contained in  $I$  in time  $O(1)$ . This matches the query times of [29].

#### 3.1 Knapsack via Convolution of Monotone Functions

First, we give a brief recap of the knapsack approach by Chan [21]. We consider the more general problem of approximating the function  $f_J: [0, t] \rightarrow \mathbb{R}_+$ , where  $J \subseteq [n]$  is a set of items and

$$f_J(x) = \max \left\{ \sum_{i \in I} p_i : \sum_{i \in I} w_i \leq x, I \subseteq J \right\}. \quad (2)$$

Intuitively, the value  $f_J(x)$  corresponds to the best possible knapsack solution if we can only pick items which are contained in  $J$  and if the weight of the solution can be *at most*  $x$ . Therefore,  $f_{[n]}(B)$  corresponds to the optimum solution of the global knapsack instance.

Note that for each  $J \subseteq [n]$ ,  $f_J(x)$  is a monotonically increasing piecewise constant function: Indeed, consider  $x' \leq x$ . Any solution  $I \subseteq J$  that is feasible for  $x'$  (i.e., the weight of  $I$  is at most  $x'$ ) is also a feasible solution for  $x$ . Thus,  $f_J(x') \leq f_J(x)$  and, therefore,  $f_J$  is monotonically increasing. Furthermore,  $f_J$  is piecewise constant since each function value  $f_J(x)$  corresponds to a solution  $I \subseteq J$  and the number of choices for  $I \subseteq J$  is finite.

Next, note that if we have two disjoint subsets  $J_1, J_2 \subseteq [n]$  then it holds that  $f_{J_1 \cup J_2}$  is the  $(\max, +)$ -convolution of  $f_{J_1}$  and  $f_{J_2}$ , i.e., for all  $x$  it holds that

$$f_{J_1 \cup J_2}(x) = \max_{\bar{x}} f_{J_1}(\bar{x}) + f_{J_2}(x - \bar{x}).$$

This can be seen by observing that for each  $x$ , the optimum solution  $I$  for the instance  $J_1 \cup J_2$  with weight at most  $x$  can be split into two disjoint solutions  $I_1 \subseteq J_1$  and  $I_2 \subseteq J_2$  such that  $I_1$  has weight  $\bar{x}$  and  $I_2$  has knapsack weight at most  $x - \bar{x}$  (for suitable choice of  $\bar{x} \in [0, x]$ ). We conclude that if we have two knapsack instances over disjoint sets of items  $J_1$  and  $J_2$ , then we compute the solution for the knapsack instance with items  $J_1 \cup J_2$  by computing the  $(\max, +)$ -convolution of  $f_{J_1}$  and  $f_{J_2}$ .

**The Exact DP.** The previous paragraphs imply a simple way of computing the exact solution of a knapsack instance: For each item  $i \in [n]$ , compute the function  $f_{\{i\}}$  and then recursively merge the solutions for sets of size  $2^j$ ,  $j = 1, \dots, \lceil \log n \rceil$ , by computing  $(\max, +)$ -convolutions until we have computed the global solution  $f_{[n]}$ . We perform the recursive merging of the solutions using a balanced binary tree, resulting in a tree of height  $O(\log n)$ .

More concretely, we build a rooted balanced binary tree  $T$  with  $n$  leaf nodes, where all edges point towards the root. We have one leaf  $f_{\{i\}}$  for each item  $i$ . Each internal node  $u$  in  $T$  is associated with a function  $f_{J_u}$  as per Equation (2), where  $J_u$  is the set of all items in the subtree rooted at  $u$ . To simplify notation, we will also refer to  $f_{J_u}$  as  $f_u$ .

Now we consider the exact computation of the DP. This will reveal the procedures  $\mathcal{P}_i$  from Definition 8. As base case, for each  $i \in [n]$ , the  $i$ 'th leaf of  $T$  contains the function  $f_{\{i\}}$ , which is a piecewise constant function that has value 0 on the interval  $[0, w_i)$  and value  $p_i$  on the interval  $[w_i, t]$ .

Next, in each internal node  $u$  of  $T$  with children  $u_1$  and  $u_2$ , we set  $f_u$  to the  $(\max, +)$ -convolution of  $f_{u_1}$  and  $f_{u_2}$ . By induction it can be seen that for every node  $u$  in  $T$ , it holds that  $J_u = J_{u_1} \cup J_{u_2}$  and thus  $J_u$  is the set of all items whose corresponding leaf is contained in the subtree  $T_u$ . Hence, for the root  $r$  of  $T$  it holds that  $f_r = f_{[n]}$  and  $f_r(B)$  is the optimal solution for the global knapsack instance.

In the following, we check that our DP satisfies Properties (1–3) of Definition 8.

First, note that the tree  $T$  from above is also the dependency graph of our DP. Hence, our DP has a row for every vertex of  $T$  and thus  $O(n)$  rows in total. Furthermore, since  $T$  has height  $O(\log n)$  and all edges point towards the root, every vertex can reach at most  $h = O(\log n)$  vertices. Hence, Property (2) of Definition 8 is satisfied.

Second, we observe that in both cases above, the function  $f_{\{i\}}$  and  $f_u$  which correspond to the rows of our DP table are monotonically increasing (we argued this above for all functions  $f_J$ ). Thus, Property (1) is satisfied.

Third, observe that Property (3) is also satisfied since  $(\max, +)$ -convolution satisfies our sensitivity condition.

We conclude that the first three properties of Definition 8 are satisfied. Unfortunately, this does not yet imply that we can obtain efficient algorithms: Note that if we compute the exact DP bottom-up then we compute one convolution per node and thus the total running time of this approach is  $O(n \cdot t(p))$ , where  $p$  is an upper bound on the number of pieces in our functions and  $t(p)$  is the time it takes to compute a  $(\max, +)$ -convolution of two functions with  $p$  pieces. However, observe that computing the convolutions can potentially take a large amount of time because the number of pieces of the functions might grow quadratically after each convolution (see Lemma 7). We will resolve this issue below using rounding.

**The Approximate DP.** Next, we consider approximations which will reveal the functions  $\tilde{\mathcal{P}}_i$  from Definition 8.

First, note that we need to compute  $(\max, +)$ -convolutions of monotonically increasing functions efficiently. We observe that this can be done efficiently using our subroutine from Lemma 7 for the  $(\min, +)$ -convolution of monotonically decreasing functions: Indeed, suppose that  $f$  is the  $(\max, +)$ -convolution of two monotonically increasing functions  $g$  and  $h$ , then for all  $x$  it holds that

$$f(x) = \max_{\bar{x}} \{g(\bar{x}) + h(x - \bar{x})\} = -\min_{\bar{x}} \{-g(\bar{x}) + (-h(x - \bar{x}))\}.$$

Now observe that  $-g$  and  $-h$  are monotonically decreasing functions and, therefore,  $f = -((-g) \oplus (-h))$ , where  $\oplus$  denotes the  $(\min, +)$ -convolution. Thus, we can use the efficient routine for  $(\min, +)$ -convolution from Lemma 7 with the same running time.<sup>6</sup>

Now we can define the subroutines  $\tilde{\mathcal{P}}_i$ . Let  $\delta > 0$  be a parameter that we set later. Whenever we compute a function  $f_u$  via a  $(\max, +)$ -convolution, we use the efficient subroutine from Lemma 7. After computing the convolution, we set  $f_u = \lceil f_u \rceil_{1+\delta}$  via the subroutine from Lemma 6.

Observe that this approach satisfies Property (4a) of Definition 8 with  $\alpha = 1 + \delta$ . Furthermore, Property (4b) is satisfied since we only use a single convolution and a single rounding step. Finally, Property (4c) is also satisfied because the resulting function is monotone and has  $p = O(\log_{1+\delta}(W))$  after the rounding.

The above arguments show that our DP is  $(h, \alpha, p)$ -well-behaved for  $h = \lceil \log n \rceil$ ,  $\alpha = 1 + \delta$ ,  $\delta = \ln(1+\epsilon)/\lceil \log n \rceil$  and  $p = O(\log_{1+\delta}(W)) = O(\log(W)/\delta)$ . Hence, Theorem 10 immediately implies the following lemma.

► **Lemma 11.** *Let  $\epsilon > 0$ . There exists an algorithm that computes a  $(1 + \epsilon)$ -approximate solution for 0-1 knapsack in time  $n \cdot \frac{1}{\epsilon^2} \log^2(n) \log^2(W) \cdot \text{polylog}(\frac{1}{\epsilon} \log(nW))$ .*

We note that we can return our solution  $I$  in time  $|I| \log(n) \cdot \text{polylog}(\frac{1}{\epsilon} \log(nW))$  as follows. Recall that our global objective function value is achieved by  $f_r(B)$  and that  $f_r(B) = f_{u_1}(\bar{x}^*) + f_{u_2}(B - \bar{x}^*)$ , where  $u_1$  and  $u_2$  are the nodes below the root node  $r$  of the dependency tree. Now using the second part of Lemma 7 we can get the value of  $\bar{x}^*$  in time  $O(\log p)$ . If  $f_{u_1}(\bar{x}^*) > 0$  we recurse on  $f_{u_1}(\bar{x}^*)$  and if  $f_{u_2}(B - \bar{x}^*) > 0$  we also recurse on  $f_{u_2}(B - \bar{x}^*)$ . At some point we will reach a leaf node  $i$  and we include  $i$  in the solution iff  $f_{\{i\}}(x) > 0$ . Note that since we only recurse for function values which are strictly larger than zero, for each item that we include into the solution we have to follow a single path in the dependency tree of height  $O(\log n)$  and our work in each internal node is bounded

<sup>6</sup> We note that, formally, Lemma 7 can only be applied on functions with non-negative values. However, this can be achieved by adding a number  $C$  to  $-g$  and  $-h$ , which is an upper bound on the values taken by  $g$  and  $h$ , and at the end we subtract the constant function  $2C$ , i.e., we set  $f = -((-g+C) \oplus (-h+C)) - 2C$ .

by  $O(\log p)$ . This gives the total time of  $O(|I| \log(n) \log(p))$  and our claim follows from our choice of  $p$  above.

**Extension to the Dynamic Setting.** Next, we extend our result to the dynamic setting. For the sake of simplicity, we assume that  $n$  is an upper bound on the maximum number of available items (items in  $S$ ) and given to our algorithm in the beginning.<sup>7</sup> We consider update operations that insert and delete items from the set. More concretely, we consider the following update operations:

- $insert(p_i, w_i)$ , in which  $i$  is added to  $S$  by setting the price and weight of item  $i$  to  $p_i \in W_\infty$  and  $w_i \in \mathbb{R}_+$ , respectively, and
- $delete(i)$ , where item  $i$  is removed from the set of items.

Our implementation is as follows. In the preprocessing phase, we build the same tree  $T$  as above and use the subroutine from above to compute the function  $f_{\{i\}}$ . For the operation  $delete(i)$ , we set  $p_i = 0$  and  $w_i = 0$ , which changes exactly one row of our DP table. For the operation  $insert(p_i, w_i)$ , we set the price and weight of item  $i$  to  $p_i$  and  $w_i$ , resp., which again changes a single row in our DP table. After changing such a row, we recompute the global DP solution via Theorem 10. Since the DP is  $(h, \alpha, p)$ -well-behaved with the same parameters as above, the theorem implies the following proposition.

► **Proposition 12.** *Let  $\epsilon > 0$ . There exists an algorithm for the fully dynamic knapsack problem that maintains a  $(1 + \epsilon)$ -approximate solution with worst-case update time  $\frac{1}{\epsilon^2} \log^3(n) \log^2(W) \cdot \text{polylog}(\frac{1}{\epsilon} \log(nW))$ .*

Observe that with the same procedure as for the static algorithm, we can return our solution  $I$  in time  $|I| \log(n) \cdot \text{polylog}(\frac{1}{\epsilon} \log(nW))$ . Furthermore, given an item  $i \in [n]$ , we can return whether  $i \in I$  in time  $\log(n) \cdot \text{polylog}(\frac{1}{\epsilon} \log(nW))$ . This can be done by using the same query procedure as in the static setting, where we only recurse on the unique subtree in the dependency tree that contains the node for item  $i$ .

We note that the above proposition already improves upon the update time in the result of Eberle et al. [29] in terms of the dependency on  $\frac{1}{\epsilon}$  but it has a worse dependency on  $\log(nW)$ . However, our query time is slower than the  $O(1)$ -time query operation in [29]. We will resolve these issues in the next subsection, where we will use the algorithm from Proposition 12 as a subroutine.

### 3.2 Dynamically Maintaining a Small Instance

Next, we obtain a faster dynamic algorithm with update time  $\tilde{O}(\frac{1}{\epsilon^2} \log^2(nW))$  by combining the algorithm from Proposition 12 and with ideas from Eberle et al. [29]. Our high-level approach is as follows. First, we partition the items into a small number of *price classes*. Then we take a few items of small weight from each price class. This will give a very small knapsack instance  $X$  for which we maintain an almost optimal solution using the subroutine from Proposition 12; since this instance is very small (i.e.,  $|X| \ll n$ ), the update time for maintaining this instance essentially becomes  $O(\frac{1}{\epsilon^2} \log^2(W))$ , i.e., we lose the  $O(\log^3 n)$  term that made the update time in the proposition too costly. For the rest of the items which are not contained in  $X$ , we show that we can compute a good solution using fractional knapsack,

<sup>7</sup> It is possible to drop this assumption using an amortization argument. More concretely, every time the number of items is less than  $n/2$  or more than  $n$ , we rebuild the data structure with a new value of  $n$ . Each rebuild can be done in time  $O(nt(n))$ , where  $t(n)$  is our update time. Since this only happens after  $\Omega(n)$  updates occurred, we can amortize this cost over the updates that appeared since the last rebuild.

which can be easily solved using a set of binary search trees. Then it remains to show that the combination of the two solutions is a  $(1 + \epsilon)$ -approximation.

The main differences of our algorithm and the one by Eberle et al. [29] are as follows. Eberle et al. also partition the items into a small number of price classes. They also combine solutions for a small set of heavy items  $X$  and solutions based on fractional knapsack for the other items. However, they have to enumerate many different sets  $X$  and they also guess the approximate price of the fractional knapsack solution; more concretely, they enumerate  $\Theta(\frac{1}{\epsilon^2} \log(W))$  choices for  $X$  and the number of guesses they have to make for the fractional knapsack solution is  $\Theta(\frac{1}{\epsilon} \log(W))$ . Thus they have to consider  $\Theta(\frac{1}{\epsilon^3} \log^2(W))$  guesses and for each of them they have to compute approximate solutions, which takes time  $\Theta(\frac{1}{\epsilon^4})$  for each  $X$  since they have to run a static algorithm from scratch. In our approach, we only have to consider a single set  $X$  which we maintain in our data structure from Proposition 12, which saves us a lot of time. Furthermore, the piecewise constant function, in which we store the solution for  $X$ , essentially “guides” our  $\Theta(\frac{1}{\epsilon} \log(W))$  guesses for the weight of fractional knapsack solution. In our analysis we have to be slightly more careful to ensure that our guesses for the weight of the fractional knapsack solution guarantee the correct approximation ratio.

**Definitions.** We assume that  $\epsilon < 1$  and that  $1/\epsilon$  is an integer. More concretely, we run the algorithm with  $\epsilon' = \max\{\frac{1}{i} : \frac{1}{i} \leq \epsilon, i \in \mathbb{N}\}$ . Set  $L = \lceil \log_{1+\epsilon}(W) \rceil$  and recall that we set  $W = \sum_i p_i$ .

We define the *price classes*  $V_\ell = \{i : (1 + \epsilon)^\ell \leq p_i < (1 + \epsilon)^{\ell+1}\}$ . In the following, we assume that all items from price class  $V_\ell$  have price exactly  $(1 + \epsilon)^{\ell+1}$ . We only lose a factor of  $1 + \epsilon$  by making this assumption. Furthermore, we set  $V_\ell^{1/\epsilon}$  to the set of  $1/\epsilon$  items from  $V_\ell$  with smallest weights  $w_i$  (breaking ties arbitrarily). We also define  $V'_\ell = V_\ell \setminus V_\ell^{1/\epsilon}$ .

Next, we set  $X = \bigcup_{\ell \geq 0} V_\ell^{1/\epsilon}$  and  $Y = \bigcup_{\ell \geq 0} V'_\ell$  for all  $\ell \geq 0$ . Note that  $X$  and  $Y$  partition the set of items and  $|X| = \frac{1}{\epsilon} \cdot L = O(\epsilon^{-2} \log(W))$ .

Now our strategy is to use our algorithm from Proposition 12 to maintain a solution for the items in  $X$ . Then we show how we can combine the solution for  $X$  with a solution for  $Y$  that is based on fractional knapsack and a charging argument.

**Data Structures.** For each  $\ell \in [L]$ , we maintain  $V_\ell$  sorted non-decreasingly by weight.

We also maintain the set  $X$  in a binary search tree, in which we sort the items by their index, and we maintain our data structure from Proposition 12 on the items in  $X$ .

Furthermore, let  $U_\ell = \bigcup_{\ell' \leq \ell} V'_{\ell'}$  denote the set of all items that are not contained in  $X$  and of price class at most  $\ell$ . For each  $\ell$ , we maintain the set  $U_\ell$  in a binary search tree  $T$  in which the items are stored as leaves and sorted by their density  $\frac{p_i}{w_i}$ . In each internal node  $u$  of  $T$ , we store the total weight of the items in the subtree  $T_u$  rooted at  $u$  and the total profit of the items in  $T_u$ . Observe that this allows us to answer queries of the type: “Given a budget  $b$ , what is the value of the optimal fractional<sup>8</sup> knapsack solution in  $U_\ell$  with weight at most  $b$ ?” in time  $O(\log n)$ .

**Updates.** Now consider an item insertion or deletion and suppose that the updated item is of price class  $V_\ell$ . We first update the sets  $V_{\ell'}$ ,  $U_{\ell'}$  for  $\ell' \leq \ell$  and the sets  $X$  and  $Y$ . Note that for each of these sets at most one item can be removed and inserted. Thus, these steps can be done in time  $O(\ell \cdot \log(n)) = O(\epsilon^{-1} \log(W) \log(n))$ .

Next, if  $X$  changed in the previous step, then we also perform the corresponding updates

<sup>8</sup> In fractional knapsack, we may use items fractionally. An optimal solution is achieved by sorting the items by their density and greedily adding items to the solution until we have used up our budget  $b$ . This approach uses at most one item fractionally (namely, the one at which we use up our budget).

in the data structure from Proposition 12. Since  $|X| = O(\epsilon^{-2} \log(W))$  holds by construction of  $X$ , the update operations for the data structure maintaining the knapsack solution for  $X$  take a total time of

$$\begin{aligned} & O\left(\epsilon^{-2} \log^3(|X|) \log^2(W) \cdot \text{polylog}\left(\frac{1}{\epsilon} \log(|X| W)\right)\right) \\ & = O\left(\epsilon^{-2} \log^2(W) \cdot \text{polylog}\left(\frac{1}{\epsilon} \log(nW)\right)\right). \end{aligned}$$

Furthermore, we can explicitly write down our solution  $I_X$  for the items in  $X$  in time  $\epsilon^{-2} \log(W) \cdot \text{polylog}(\frac{1}{\epsilon} \log(nW))$  since  $|X| = O(\epsilon^{-2} \log(W))$ . Also, for each  $i \in I_X$ , we can set a bit indicating that  $i \in I_X$ . Note that the time for writing down  $I_X$  and setting the bits is subsumed by the update time above.

**Queries.** *Returning the value of a solution:* We return the value of a global knapsack solution as follows.

Consider the data structure from Proposition 12 which maintains a solution for the items in  $X$ . Note that this solution is stored as a piecewise constant function with  $p \leq L$  pieces and consider the list representation  $(x_1, y_1), \dots, (x_p, y_p)$  of this function.

Our strategy is as follows: For each  $i = 1, \dots, p$ , we consider a solution which spends budget  $x_i$  on items in  $X$  and budget  $B - x_i$  on items in  $Y$ . Then we take the maximum over all of the solutions we have considered. More concretely, for given  $i = 1, \dots, p$ , we obtain our solution as follows. We pick  $\ell_i$  such that  $(1 + \epsilon)^{\ell_i} = \lceil \epsilon \cdot y_i \rceil_{1+\epsilon}$  (see Lemma 13 below for a justification of this choice). Now we use the binary search tree for  $U_{\ell_i}$  to find the highest profit that we can obtain from fractional knapsack on items in  $U_{\ell_i} \subseteq Y$  if we can spend budget at most  $b = B - x_i$ . Let  $y'_i$  be the value of this query after removing any profit that we gain from the (at most one) fractionally cut item. We also store the density of the final item that is contained in the fractional knapsack solution. Now we return the maximum of  $y_i + y'_i$  over all  $i = 1, \dots, p$ .

Note that since the solution for  $X$  has at most  $L = O(\epsilon^{-1} \log(W))$  pieces and for each of them we perform a single query in a binary search tree, the total time for return the solution value is  $O(\epsilon^{-1} \log(W) \log(n))$ . Note that this time is subsumed by the update time.

*Returning the entire solution:* Now we can return our global solution  $I$  in time  $O(|I|)$  as follows. Observe that  $I$  is composed of the solution  $I_X$  for the items in  $X$  and of the items in the fractional knapsack solution. During our updates, we already stored the items in  $I_X$  and can write them down in time  $O(|I_X|)$ . Next, to return the items from the fractional knapsack solution, recall that we stored the density of the final item in the fractional knapsack solution. Thus, we only have to output the items ordered non-decreasingly by their density, while we are above the desired density-threshold. This can be done in time linear in the size of the fractional knapsack solution. This is essentially the same query procedure as in [29].

*Returning whether an item is in the solution:* Furthermore, observe that the above implies that we can answer whether an item  $i \in [n]$  is contained in our solution in time  $O(1)$ : If  $i \in X$  then we already stored a bit whether  $i \in I_X$ . If  $i \notin X$  then we can check whether  $i$  is in the fractional knapsack solution by checking whether its density is above or below the threshold given by the final item in the fractional knapsack solution.

**Analysis.** We start by making some simplifications to OPT. We let  $\text{OPT}'$  denote the version of OPT in which for each  $\ell \in [L]$ , we pick the  $|\text{OPT}' \cap V_\ell|$  items of smallest weight from  $V_\ell$ . This only loses a factor of  $1 + \epsilon$ . Next, define  $\text{OPT}'_X = \text{OPT}' \cap X$  and  $\text{OPT}'_Y = \text{OPT}' \cap Y$ . Observe that by how we picked  $\text{OPT}'$ , it holds that  $\text{OPT}'_Y \cap V_\ell \neq \emptyset$  iff  $|\text{OPT}' \cap V_\ell| > 1/\epsilon$ .

Let  $p_X$  denote the total price of items in  $\text{OPT}'_X$  and let  $w_X$  denote the total weight of the items in  $\text{OPT}'_X$ . Let  $f$  denote the piecewise constant function that stores the solution



for the items in  $X$ . Observe that by Proposition 12 we have that

$$p_X \leq f(w_X) \leq (1 + \epsilon)p_X.$$

Also, the function value  $f(w_X)$  is part of a piece  $(x_{i^*}, y_{i^*})$  with  $x_{i^*} \leq w_X$  and  $y_{i^*} = f(w_X)$ .

The next lemma justifies why we set  $\ell_i$  such that  $(1 + \epsilon)^{\ell_i} = \lceil \epsilon \cdot y_i \rceil_{1+\epsilon}$  in our algorithm. To this end, let  $\ell_{i^*}$  be such that  $(1 + \epsilon)^{\ell_{i^*}} = \lceil \epsilon \cdot y_{i^*} \rceil_{1+\epsilon}$  and let  $\ell_Y$  be the price class of the most valuable item in  $\text{OPT}'_Y$ . In the lemma we show that  $\ell_{i^*} \geq \ell_Y$ . We will use this to show that our solution for  $X$  of profit  $y_{i^*}$  is valuable enough such that we can charge a fractionally cut item from fractional knapsack onto the solution from  $X$  and only lose a factor of  $(1 + \epsilon)^2$ .

► **Lemma 13.** *It holds that  $\ell_{i^*} \geq \ell_Y$ .*

**Proof.** Since  $\text{OPT}'_Y \cap V'_{\ell_Y} \neq \emptyset$ ,  $|\text{OPT}' \cap V_{\ell_Y}| > 1/\epsilon$  and thus  $\text{OPT}'_X$  contains all  $1/\epsilon$  items from  $V_{\ell_Y}^{1/\epsilon}$ . Hence,  $p_X \geq \frac{1}{\epsilon} \cdot (1 + \epsilon)^{\ell_Y}$ . From above we get  $f(w_X) = y_{i^*}$  and  $f(w_X) \geq p_X$ . By choice of  $\ell_{i^*}$ ,

$$(1 + \epsilon)^{\ell_{i^*}} = \lceil \epsilon \cdot y_{i^*} \rceil_{1+\epsilon} = \lceil \epsilon \cdot f(w_X) \rceil_{1+\epsilon} \geq \lceil \epsilon \cdot p_X \rceil_{1+\epsilon} \geq \left\lceil \epsilon \cdot \frac{1}{\epsilon} (1 + \epsilon)^{\ell_Y} \right\rceil_{1+\epsilon} = (1 + \epsilon)^{\ell_Y}.$$

This implies  $\ell_{i^*} \geq \ell_Y$ . ◀

Next, consider the the fractional knapsack solution that we obtain from our query. Note that this solution has a profit that is at least as large as the profit of  $\text{OPT}'_Y$  (since fractional knapsack is a relaxation of 0-1 knapsack). Furthermore, the fractional solution uses at most one item fractionally and this item is from  $U_{\ell_{i^*}}$  and has value at most  $(1 + \epsilon)^{\ell_{i^*}} = \lceil \epsilon \cdot y_{i^*} \rceil_{1+\epsilon} \leq (1 + \epsilon)\epsilon \cdot y_{i^*}$ . Thus, we can charge this item on  $\text{OPT}'_X$  and lose a factor of at most  $(1 + \epsilon)^2$ .

We conclude that the solution  $y_{i^*} + y'_{i^*}$  is a  $(1 + \epsilon)^{O(1)}$ -approximation of  $\text{OPT}$ . Combining this with our previous running time analysis, we obtain Theorem 1.

## 4 Technical Overview

We now present an overview of two techniques for making DPs fit our framework. We will briefly discuss how we *monotonized* the DP for  $k$ -balanced partitioning and how we *inverted* the DP for simultaneous source location. Due to space constraints, we only present excerpts of our algorithms and we only consider special cases. More concretely, for both problems we will consider the special case when the input graph is a binary tree. In the appendix we will show that the results can be extended to general graphs.

### 4.1 Monotonizing the DP of Feldmann and Foschini

We start by considering the  $k$ -balanced graph partitioning problem. Recall that in this problem, the input is a graph  $G = (V, E, \text{cap})$ , where  $\text{cap} : E \rightarrow W_\infty$  is a weight function on the edges, and an integer  $k$ . As discussed in the introduction, we assume that we can violate the partition sizes by a  $(1 + \epsilon)$ -factor and our goal is to find a partition  $V_1, \dots, V_k$  of the vertices such that  $|V_i| \leq \lceil (1 + \epsilon)|V|/k \rceil$  for all  $i$  and such that we minimize  $\text{cut}(V_1, \dots, V_k) := \sum_{i=1}^k \sum_{\{u,v\} \in E \cap (V_i \times (V \setminus V_i))} \text{cap}(u, v)$ .

For the sake of better exposition, here we only consider the special case in which  $G$  is a *binary tree*; in Appendix D.4 we show how to drop this assumption.

In the following we present a DP in which the rows are monotone and we show how to efficiently perform operations on these solution vectors using monotone piecewise constant functions. Our DP is related to the DP by Feldmann and Foschini [34] which is *non-monotone* and thus our DP can be viewed as the *monotonization* of the DP by Feldmann and Foschini. We believe that our technique to monotonize the DP will have further applications in the future.

**High-Level Description of the DP.** Our DP is computed bottom-up starting at the leaves of the tree and then moving up in the tree. For each vertex  $v$ , we will compute a DP solution of minimum cost that encodes whether the edge to the parent  $p$  of  $v$  is cut and which edges shall be cut inside the subtree  $T_v$  that is rooted at  $v$ . Note that the removal of the cut edges in our solution will decompose the tree into disjoint connected components and exactly one of them contains  $v$ 's parent  $p$ . Additionally, we store information about the number of vertices that are still connected to the parent  $p$  (and, therefore, to the outside of  $T_v$ ) after the cut edges are removed. We will assume that when we compute the DP cell for a vertex  $v$ , we have access to the solutions for both of its children.

More concretely, when we have computed a solution for a subtree  $T_v$ , i.e., we know which edges incident to nodes in this subtree we are going to remove (note that the edge leading to the parent of  $v$  is incident to  $T_v$  and thus we consider it as part of this solution), we store the following information in the DP table. First, we store its cost, i.e., the total capacity of all edges that are incident to vertices in  $T_v$  and that are cut. As described above, we would also like to store the number of vertices that are connected to the parent of  $v$  and the sizes of connected components inside  $T_v$ . However, there are two difficulties: (1) We cannot store the number of vertices that are connected to the root exactly because this would result in a too large DP table. Instead, we store the cheapest solution in which vertices of *at most* some given number are still connected to the parent of  $v$ . As we will see, this approach gives rise to *monotonically decreasing* functions and allows for a very efficient computation of the DP table. (2) We store implicitly the *size* of all connected components that are created after the cut edges are removed and that lie completely inside  $T_v$ . As before, storing these sizes exactly would result in a very large DP table and, therefore, we store them concisely using the concept of a *signature*. The signatures will help us to characterize the sizes of the components inside  $T_v$  very efficiently.

**Signatures.** We call a connected component in  $T_v$  *large* if it contains at least  $\epsilon \lceil |V|/k \rceil$  vertices and otherwise we call it *small*. Let  $t = \lceil \log_{1+\epsilon}(1/\epsilon) \rceil + 1$ , and let  $M = \lceil k/\epsilon \rceil + 1$ . A *signature* is a vector  $g = (g_0, \dots, g_{t-1}) \in [M-1]^t$ . Observe that each  $\mathcal{P}_i$  is an integer between 0 and  $M-1$  and hence there are  $M^t = (k/\epsilon)^{O(\epsilon^{-1} \log(1/\epsilon))}$  different signatures. Intuitively, an entry  $\mathcal{P}_i$  in  $g$  tells us roughly how many components of size  $(1+\epsilon)^i \cdot \epsilon \lceil |V|/k \rceil$  there are in the DP solutions that we consider. Due to space constraints, we refer to the appendix for the formal definition.

For  $x \in \mathbb{N}$ , we let  $e(x) \in [M-1]^t$  denote the signature of a single component with  $x$  vertices. More precisely, we set  $e(x)$  to the vector that has  $e(x)_j = 1$  for  $j = \arg \min\{j \in \mathbb{N} : x \leq (1+\epsilon)^j \cdot \epsilon \lceil |V|/k \rceil\}$  and  $e(x)_j = 0$ , otherwise. If  $x < \epsilon \lceil |V|/k \rceil$ , we define  $e(x) = \vec{0}$ .

**Formal DP Definition.** Now we describe the DP formally. An entry  $\text{DP}(v, g, \text{cut}, x) \in W_\infty$  in the DP table for a vertex  $v$  is indexed by a signature  $g$ , a Boolean value  $\text{cut}$  and  $x \in [n]$ . We will consider the tuples  $(v, g, \text{cut})$  as the rows  $\mathcal{I}$  of the DP table and  $x$  as the columns; we associate each such row with a function  $\text{DP}(v, g, \text{cut}, \cdot) : [n] \rightarrow W_\infty$ . Note that our DP has  $|V| \cdot M^t \cdot 2 = (k/\epsilon)^{O(\epsilon^{-1} \log(1/\epsilon))} \cdot n$  rows. Also, note that it has columns  $n$ ; later, even though  $x$  only takes discrete values, we will allow  $x$  to take values in  $[0, \infty)$ .

An entry  $\text{DP}(v, g, \text{cut}, x)$  describes the optimum cost of cutting edges incident on the

subtree  $T_v$  (including the cost of maybe cutting the edge to the parent of  $v$ ). We will refer to the set of vertices in  $T_v$  that are still connected to the parent of  $v$  after the cut edges are removed as the *root component*. We impose the following conditions on  $\text{DP}(v, g, \text{cut}, x)$ :

- Once the cut edges are removed, the root component  $U \subseteq T_v$  has *at most*  $x$  vertices, i.e.,  $|U| \leq x$ .
- If *cut* is set to true then the edge between  $v$  and its parent is cut, otherwise it is kept.
- The vertices inside  $T_v$  that (once the cut edges are removed) are *not* connected to the parent of  $v$  form connected components that are consistent with the signature  $g$ .

Next, we observe that if we fix a vertex  $v$ , a signature  $g$  and a value for *cut*, then the resulting function  $\text{DP}(v, g, \text{cut}, \cdot)$  is monotonically decreasing in  $x$ .

► **Observation 14.** *Let  $v \in V$ ,  $g \in [M - 1]^t$  be a signature and  $\text{cut} \in \{\text{true}, \text{false}\}$ . Then the function  $\text{DP}(v, g, \text{cut}, \cdot) : [0, \infty) \rightarrow \mathbb{R}_+$  is monotonically decreasing.*

**Proof.** By definition,  $\text{DP}(v, g, \text{cut}, x)$  stores the cost of the optimum solution in which there are *at most*  $x$  vertices in the root component. Since  $x \leq x'$ , the solution  $\text{DP}(v, g, \text{cut}, x)$  is also a feasible solution for  $\text{DP}(v, g, \text{cut}, x')$ . Hence,  $\text{DP}(v, g, \text{cut}, \cdot)$  is monotonically decreasing. ◀

**Comparison With the DP by Feldmann and Foschini.** When comparing our DP with the one by Feldmann and Foschini [34] then one of the crucial changes is that in our DP,  $x$  encodes an *upper bound* on the number of vertices in the root component. Previously, Feldmann and Foschini considered root components with *exactly*  $x$  vertices. This is why their DP was *non-monotone* and why one can view our DP as the *monotonization* of the DP in [34]. However, we also generalize the DP to the setting with vertex weights and, as we will see below, parts of our algorithm for computing the DP approximately are rather involved.

### 4.1.1 Computing the DP

We now give a flavor of what our algorithms for computing the DP look like. We start by showing how to compute the exact DP solution  $\text{DP}(v, \cdot, \cdot, \cdot)$  for a vertex  $v$  of the tree, where  $v$  has parent  $p$  and children  $v_l, v_r$  and it is connected to them via edges  $e_p, e_l$  and  $e_r$ , respectively.

Computing the DP is based on several case distinctions; here, we only consider the case *in which  $v$  is an internal vertex of the tree we do not cut the edges  $e_l$  and  $e_r$* . All other cases are presented in the appendix.

When computing a DP row given by  $\text{DP}(v, \cdot, \cdot, \cdot)$ , we will only require access to the DP rows  $\text{DP}(v_l, \cdot, \cdot, \cdot)$  and  $\text{DP}(v_r, \cdot, \cdot, \cdot)$ . This implies that the dependency tree of the DP is a tree and has the same height as our input graph  $G$  (recall that here we assume that  $G$  is a binary tree). Note that the height of the tree also implies an upper bound on the number of reachable nodes.

**Exact Computation.** We start with the exact computation. Here, we can afford to iterate over all values  $x \in [n]$  and  $g \in [M - 1]^t$  to compute  $\text{DP}(v, \cdot, \cdot, \cdot)$ . Therefore, we consider the values for  $x$  and  $g$  as input to our algorithm.

Since we assume that we do not cut the edges  $e_l$  and  $e_r$ , we have to select subsolutions for  $T_{v_l}$  and  $T_{v_r}$ , where each subsolution is characterized by the upper bound  $x_l$  (resp.  $x_r$ ) and its signature  $g_l$  (resp.  $g_r$ ).

First, suppose that we cut the edge  $e_p$ . If we let  $x_l$  and  $x_r$  denote the number of vertices of the root components for the subsolutions, then the vertex  $v$  will be included in a component of size  $x_l + x_r + 1$  afterwards. Hence, we can combine the subsolutions to a solution for

signature  $g$  as long as  $g_l + g_r + e(x_l + x_r + 1) = g$ . Consequently we set for every  $x \in [0, \infty)$ ,

$$\text{DP}_B(v, g, \text{true}, x) = \text{cap}(v, p) + \min_{x_l, x_r, g_l + g_r = g - e(x_l + x_r + 1)} \text{DP}(v_l, g_l, \text{false}, x_l) + \text{DP}(v_r, g_r, \text{false}, x_r).$$

Second, suppose that we do not cut  $e_p$ . Again we have to set  $\text{DP}_B(v, g, \text{false}, x) = \infty$  for all signatures  $g$  and all  $x \in [0, 1)$ , because  $v$  can reach  $p$ . For  $x \geq 1$  we have to select  $x_l$  and  $x_r$  such that they sum to  $x - 1$  as this guarantees that at most  $x$  vertices can reach the parent  $p$ . Consequently, we set for all  $x \in [1, \infty)$

$$\text{DP}_B(v, g, \text{false}, x) = \min_{g_l + g_r = g, x_l + x_r = x - 1} \text{DP}(v_l, g_l, \text{false}, x_l) + \text{DP}(v_r, g_r, \text{false}, x_r).$$

Here, we can afford to exhaustively enumerate all  $O(M^t n^2)$  possibilities in the min-operations above.

**Approximate Computation.** Now let us consider the approximate computation. We denote the approximate DP solution by ADP. We assume that we have already computed the children solutions  $\text{ADP}(v_l, g, \text{cut}, \cdot)$  and  $\text{ADP}(v_r, g, \text{cut}, \cdot)$  and that they are stored using our data structure from Section 2. We will maintain as an invariant that each of these functions has at most  $p = O(\log_{1+\delta}(W))$  pieces and we will ensure this *by rounding our solution* at the end of every step, i.e., by setting  $\text{ADP}(v, g, \text{cut}, \cdot) = \lceil \text{ADP}(v, g, \text{cut}, \cdot) \rceil_{1+\delta}$  using the rounding procedure from Lemma 6. This will ensure the following two properties: (1) The functions  $\text{ADP}(v, g, \text{cut}, \cdot)$  never have more than  $O(\log_{1+\delta}(W))$  pieces by Lemma 6. Thus, we can perform all of our operations very efficiently. (2) For the function at the root of the tree, the approximation error is at most  $(1 + \delta)^h$ , where  $h$  is the height of the tree. By picking  $\delta = O(\epsilon/h)$ , we will achieve that we obtain a  $(1 + \epsilon)$ -approximate solution at the root. Now we proceed to the explanation of our computation.

If we do not cut the edge to the parent of  $v$ , we proceed similar to the exact DP above. We start by setting  $\text{ADP}_B(v, g, \text{false}, x) = \infty$  for all  $x \in [0, 1)$ . Next, for  $x \in [1, \infty)$  we wish to set

$$\text{ADP}_B(v, g, \text{false}, x) = \min_{g_l + g_r = g, x_l + x_r = x - 1} \text{ADP}(v_l, g_l, \text{false}, x_l) + \text{ADP}(v_r, g_r, \text{false}, x_r) \quad (3)$$

$$= \min_{g_l + g_r = g} \min_{x_l + x_r = x - 1} \text{ADP}(v_l, g_l, \text{false}, x_l) + \text{ADP}(v_r, g_r, \text{false}, x_r). \quad (4)$$

Note that for fixed  $g_l$  and  $g_r$ , the inner min-operation in the second line describes a  $(\min, +)$ -convolution due to the constraint  $x_l + x_r = x - 1$ . Therefore, in the inner min-operation we compute a convolution  $\text{ADP}(v_l, g_l, \text{false}, \cdot) \oplus \text{ADP}(v_r, g_r, \text{false}, \cdot)$  and shift the result by 1 via the shift operation from Lemma 6 (where for  $x \in [0, 1)$  we set  $\text{ADP}_B(v, g, \text{false}, x) = \infty$ ). We need time  $O(p^2 \log p)$  for computing the convolution according to Lemma 7. To compute the outer minimum in Equation (4), we iterate over all  $g_l \in [M - 1]^t$  using Lemma 6 and thus perform  $O(M^t)$  minimum computations over piecewise constant functions with at most  $p^2$  pieces. Hence, we need time  $O(M^t p^2 \log(M^t p^2))$  according to Lemma 21. By Lemma 22,  $\text{ADP}_B(v, g, \text{false}, \cdot)$  is monotonically decreasing since it is the minimum over convolutions of two monotonically decreasing functions.

If we cut the edge to the parent of  $v$ , then for all  $x \in [0, \infty)$  we would like to set

$$\text{ADP}_B(v, g, \text{true}, x) = \text{cap}(v, p) + \min_{x_l, x_r, g_l + g_r = g - e(x_l + x_r + 1)} \text{ADP}(v_l, g_l, \text{false}, x_l) + \text{ADP}(v_r, g_r, \text{false}, x_r).$$

Note that here we need to be careful as the range of  $g_l$  and  $g_r$  depends on the choice of  $x_l + x_r$ . Since there are  $\Omega(n)$  possible values for  $x_l + x_r$ , we cannot afford to iterate over all values that

$x_l + x_r$  can take. Instead, we will show that we only need to consider  $O(\log(k/\epsilon)/\epsilon)$  different pairs  $(x_l, x_r)$  by exploiting the monotonicity of  $\text{ADP}(v_l, g_l, \text{false}, \cdot)$  and  $\text{ADP}(v_r, g_r, \text{false}, \cdot)$ .

First, observe that we can assume  $x_l \leq |T_{v_l}|$  and  $x_r \leq |T_{v_r}|$ : increasing the upper bounds on the number of vertices of the root component further would mean that the root component contains than *all* vertices inside the sub-tree, which is impossible. Thus,  $x_l + x_r + 1 \in [1, n]$ .

Second, we partition the interval  $[1, n]$  into  $O(\log(k/\epsilon)/\epsilon)$  intervals. We have intervals  $I_j = (\xi_{j-1}, \xi_j]$  with  $\xi_j = (1 + \epsilon)^j \lceil n/k \rceil$  for all  $j = 1, \dots, \log_{1+\epsilon}(k/\epsilon)$ . In addition, we add an “interval”  $I_0 := [\epsilon \lceil n/k \rceil, \epsilon \lceil n/k \rceil]$  and the interval  $I_{-1} := [1, \epsilon \lceil n/k \rceil)$ . We set  $\xi_0 = \epsilon \lceil n/k \rceil$  and we set  $\xi_{-1}$  to the largest integer that is less than  $\epsilon \lceil n/k \rceil$ . Observe that for all  $j \geq -1$  and  $x \in I_j$ , we have  $e(x) = e(\xi_j)$ , i.e., the value of  $e(x)$  does not change inside the interval  $I_j$ . Below, this property will allow us to separate the conditions on  $x_l + x_r$  and on  $g_l + g_r$ .

Now we can rewrite the above expression as

$$\begin{aligned} \text{ADP}_B(v, g, \text{true}, x) = \\ \text{cap}(v, p) + \min_j \min_{x_l + x_r + 1 \in I_j} \min_{g_l + g_r = g - e(\xi_j)} \text{ADP}(v_l, g_l, \text{false}, x_l) + \text{ADP}(v_r, g_r, \text{false}, x_r). \end{aligned}$$

Third, note that now the two min-operations only depend on the choice of  $j$  and, importantly, the minimum over  $g_l$  and  $g_r$  does not depend on the choice of  $x_l + x_r$  anymore. Therefore, we can swap the order of the two min-operations. Furthermore, since  $\text{ADP}_B(v, g, \text{false}, x)$  is monotonically decreasing with  $x$ , we can restrict the choice of  $x_l$  and  $x_r$  such that  $x_l + x_r + 1$  is the largest number in the corresponding interval  $I_j$ , i.e.,  $x_l + x_r + 1 = \xi_j$ . Thus,

$$\begin{aligned} \text{ADP}_B(v, g, \text{true}, x) = \\ \text{cap}(v, p) + \min_j \min_{g_l + g_r = g - e(\xi_j)} \min_{x_l + x_r + 1 = \xi_j} \text{ADP}(v_l, g_l, \text{false}, x_l) + \text{ADP}(v_r, g_r, \text{false}, \xi_j - x_l - 1). \end{aligned}$$

Next, we explain how the above expression can be computed efficiently. Let us first argue how we can efficiently compute the inner min-operation of the above expression. We start by observing that this min-operation is *not* a convolution since in the constraint we sum up to  $\xi_i$  which is a constant (rather than to the variable  $x$ ). Now recall that  $\text{ADP}(v_l, g_l, \text{false}, \cdot)$  and  $\text{ADP}(v_r, g_r, \text{false}, \cdot)$  are piecewise constant functions with  $O(p)$  pieces by our invariants. Since  $x_l, x_r \geq 0$  this implies that there are only  $O(p^2)$  choices for  $x_l$  and  $x_r$  such that  $x_l, x_r \in I_j$  and either a new piece starts in  $\text{ADP}(v_l, g_l, \text{false}, x_l)$  or in  $\text{ADP}(v_r, g_r, \text{false}, x_r)$ . Thus, we can iterate over all these pairs  $(x_l, x_r)$  and evaluate  $\text{ADP}(v_l, g_l, \text{false}, x_l) + \text{ADP}(v_r, g_r, \text{false}, x_r)$ , where  $x_r = \xi_j - x_l - 1$ . Thus, we can compute the inner min-operation in time  $O(p^2 \log p)$ . We note that since this min-operation is considering a super-constant number of terms, this DP is not well-behaved (it violates Property (4b) of Definition 8). This is why in our analysis we will use the more general notion from Section C.2.

Next, we can compute the outer two min-operations by simply iterating over  $j$  and all choices for  $g_l$  and setting  $g_r = g - e(\xi_j) - g_l$  as above in  $O(M^t \cdot \log(k/\epsilon)/\epsilon)$  iterations. Hence, we obtain a running time of  $O(M^t p^2 \log p \cdot \log(k/\epsilon)/\epsilon)$ .

Finally, we note that as  $\text{ADP}_B(v, g, \text{true}, x)$  is independent of  $x$ , it is a constant. Thus,  $\text{ADP}_B(v, g, \text{true}, x)$  is a piecewise constant function with a single piece and it is monotonically decreasing.

*Rounding Step.* As noted earlier, after computing the solutions  $\text{ADP}_B(v, g, \text{false}, \cdot)$  and  $\text{ADP}_B(v, g, \text{true}, \cdot)$ , we also round the solution by setting  $\text{ADP}_B(v, g, \text{cut}, \cdot) = \lceil \text{ADP}_B(v, g, \text{cut}, \cdot) \rceil_{1+\delta}$  for  $\text{cut} \in \{\text{true}, \text{false}\}$  to ensure that we only have  $p = O(\log_{1+\delta}(W))$  pieces in the resulting function. Note that this is the only approximate operation we perform and all other operations above have been exact.

## 4.2 Inverting the DP of Andreev et al.

Now we briefly describe our DP for simultaneous source location. Recall that in this problem, the input consists of an undirected graph  $G = (V, E, \text{cap}, d)$  with a capacity function  $\text{cap}: E \rightarrow W_\infty$  and a demand function  $d: V \rightarrow W_\infty$ . The goal is to select a minimum set  $S \subseteq V$  of *sources* that can simultaneously supply all vertex demands. More concretely, a set of sources  $S$  is *feasible* if there exists a flow from the vertices in  $S$  that supplies demand  $d(v)$  to all vertices  $v \in V$  and that does not violate the capacity constraints on the edges. The objective is to find a feasible set of sources of minimum size.

Here, we will again assume the special case in which  $G$  is a *binary tree*; we show in Appendix E.3 how to drop this assumption.

**DP Definition.** Given a vertex  $v$  and a value  $x \in \mathbb{R}$ , we let  $\text{DP}(v, x)$  denote the minimum number of sources that we need to place in the subtree  $T_v$  such that when  $v$  receives flow at most  $x$  from its parent then all demands in  $T_v$  can be satisfied. We note that  $x$  can take positive and negative values: for  $x \geq 0$  this corresponds to the setting in which flow is sent from the parent of  $v$  into  $T_v$  and for  $x < 0$  this corresponds to the setting in which flow is sent from  $T_v$  towards the parent of  $v$ . We further follow the convention that when the demands in  $T_v$  cannot be satisfied when  $v$  receives flow  $x$  from its parent, then we set  $\text{DP}(v, x) = \infty$ .

Observe that this DP has rows  $\mathcal{I} = V$  and columns  $\mathcal{J} = \mathbb{R}$ . Furthermore,  $\text{DP}(v, \cdot)$  is monotonically decreasing since for  $x < x'$ , any solution in which  $T_v$  receives flow at most  $x$  from the parent of  $v$  is also feasible when  $T_v$  receives flow at most  $x'$  from the parent of  $v$ . This satisfies Property (1) of Definition 8.

**The Inverse DP.** Interestingly, our DP is very related to the one by Andreev et al. [4]. They defined a function  $f(v, i)$  which, given a vertex  $v$  and an integer  $i \in \mathbb{N}$ , denotes the *minimum amount of flow* that  $v$  needs to receive from its parent if all demands in  $T_v$  need to be satisfied and if we can place  $i$  sources in the subtree  $T_v$ . Similar to above,  $f(v, i)$  takes positive values if the demand in  $T_v$  can only be satisfied by receiving flow from the parent of  $v$  and it takes negative values if the demand in  $T_v$  is already satisfied by the sources in the subtree  $T_v$  and  $v$  can send flow to its parent.

Now observe that our DP can essentially be viewed as the “inverse” of  $f(v, i)$ . More formally, observe that  $\text{DP}(v, x) = f^{-1}(v, x) := \min\{i: f(v, i) \leq x\}$ .

The reason why we chose the inverse formulation for our DP is as follows. To ensure that our algorithms are efficient, we have to make sure that our monotone piecewise constant functions have only few pieces. One natural way to do is using rounding. However, since the function values of  $f$  are positive and negative, it is not clear how we should perform the rounding. For example, to only use a small number of pieces for representing  $f$ , we would have to use different rounding mechanisms for those function values in  $[-1, 1]$  and those in  $[-W, W] \setminus [-1, 1]$ , where  $W$  is the largest edge capacity: Indeed, if we rounded the values of  $f$  to powers of  $(1 + \delta)^j$  then there are only  $O(\log_{1+\delta}(W))$  function values in  $[-W, W] \setminus [-1, 1]$  but there are infinitely many function values in  $[-1, 1]$ . Similarly, if we rounded to multiples of  $\delta$  then there are only  $O(1/\delta)$  function values in  $[-1, 1]$  but this would lead to  $O(W/\delta)$  function values in  $[-W, W] \setminus [-1, 1]$ . In both cases, our functions would have too many pieces and we would have to pick a rounding function which provides a tradeoff between these two cases. Furthermore, we would have to find an analysis that shows that this “more involved” rounding function does not introduce much too error.

In our DP we bypass these issues because we move the negative numbers into the domain of the function  $\text{DP}(v, \cdot): \mathbb{R} \rightarrow [n + 1]$ . Then in the codomain we only have non-negative numbers to which we can apply the standard rounding function  $\lceil \cdot \rceil_{1+\delta}$  in a straightforward way. This also has the positive side effects that instead of getting factors of  $\text{polylog}(W)$  in

our running times, we only get factors of  $\text{polylog}(n)$  because our codomain became  $[n + 1]$  rather than some potentially large interval  $[-W, W]$ . We believe this technique of considering inverse DPs will be useful in the future to compute approximate solutions for DPs that can take positive *and negative* values.

Due to lack of space, we present the details for computing the DP in Appendix E.

---

## References

---

- 1 Amir Abboud, Vincent Cohen-Addad, and Philip N Klein. New hardness results for planar graph problems in  $p$  and an algorithm for sparsest cut. In *STOC*, pages 996–1009, 2020.
- 2 Alok Aggarwal, Maria M. Klawe, Shlomo Moran, Peter W. Shor, and Robert E. Wilber. Geometric applications of a matrix-searching algorithm. *Algorithmica*, 2:195–208, 1987.
- 3 Greg Aloupis, Thomas Fevens, Stefan Langerman, Tomomi Matsui, Antonio Mesa, Yurai Nuñez, David Rappaport, and Godfried Toussaint. Algorithms for computing geometric measures of melodic similarity. *Computer Music Journal*, 30(3):67–76, 2006.
- 4 Konstantin Andreev, Charles Garrod, Daniel Golovin, Bruce M. Maggs, and Adam Meyerson. Simultaneous source location. *ACM Trans. Algorithms*, 6(1):16:1–16:17, 2009.
- 5 Konstantin Andreev and Harald Räcke. Balanced graph partitioning. *Theory of Computing Systems*, 39(6):929–939, 2006.
- 6 Ali Aouad and Danny Segev. An approximate dynamic programming approach to the incremental knapsack problem. *Operations Research*, 2022.
- 7 Kouji Arata, Satoru Iwata, Kazuhisa Makino, and Satoru Fujishige. Locating sources to meet flow demands in undirected networks. *Journal of Algorithms*, 42(1):54–68, 2002.
- 8 Yoan José Pinzón Ardila, Raphaël Clifford, and Manal Mohamed. Necklace swap problem for rhythmic similarity measures. In *International Symposium on String Processing and Information Retrieval*, pages 234–245. Springer, 2005.
- 9 Kyriakos Axiotis and Christos Tzamos. Capacitated dynamic programming: Faster knapsack and graph algorithms. In *ICALP*, pages 19:1–19:13, 2019.
- 10 Arturs Backurs, Piotr Indyk, and Ludwig Schmidt. Better approximations for tree sparsity in nearly-linear time. In *SODA*, pages 2215–2229, 2017.
- 11 MohammadHossein Bateni, MohammadTaghi Hajiaghayi, Saeed Seddighin, and Cliff Stein. Fast algorithms for knapsack via convolution and prediction. In *STOC*, pages 1269–1282, 2018.
- 12 Wolfgang W. Bein, Mordecai J. Golin, Lawrence L. Larmore, and Yan Zhang. The knuth–yao quadrangle-inequality speedup is a consequence of total monotonicity. *ACM Trans. Algorithms*, 6(1):17:1–17:22, 2009.
- 13 Marcin Bienkowski, Mirosław Korzeniowski, and Harald Räcke. A practical algorithm for constructing oblivious routing schemes. In *SPAA*, pages 24–33, 2003.
- 14 David Bremner, Timothy M. Chan, Erik D. Demaine, Jeff Erickson, Ferran Hurtado, John Iacono, Stefan Langerman, Mihai Pătraşcu, and Perouz Taslakian. Necklaces, convolutions, and  $X+Y$ . *Algorithmica*, 69(2):294–314, 2014.
- 15 Karl Bringmann and Alejandro Cassis. Faster knapsack algorithms via bounded monotone min-plus-convolution. In *ICALP*, volume 229, pages 31:1–31:21, 2022.
- 16 Karl Bringmann, Fabrizio Grandoni, Barna Saha, and Virginia Vassilevska Williams. Truly subcubic algorithms for language edit distance and RNA folding via fast bounded-difference min-plus product. *SIAM J. Comput.*, 48(2):481–512, 2019.
- 17 Karl Bringmann, Marvin Künnemann, and Karol Wegrzycki. Approximating APSP without scaling: equivalence of approximate min-plus and exact min-max. In *STOC*, pages 943–954. ACM, 2019.
- 18 Aydin Buluç, Henning Meyerhenke, Ilya Safro, Peter Sanders, and Christian Schulz. Recent advances in graph partitioning. In *Algorithm Engineering - Selected Results and Surveys*, volume 9220 of *Lecture Notes in Computer Science*, pages 117–158. 2016.

- 19 Rainer E. Burkard, Bettina Klinz, and Rüdiger Rudolf. Perspectives of monge properties in optimization. *Discret. Appl. Math.*, 70(2):95–161, 1996.
- 20 Michael Bussieck, Hannes Hassler, Gerhard J Woeginger, and Uwe T Zimmermann. Fast algorithms for the maximum convolution problem. *Operations research letters*, 15(3):133–141, 1994.
- 21 Timothy M. Chan. Approximation schemes for 0-1 knapsack. In *SOSA*, volume 61, pages 5:1–5:12, 2018.
- 22 Timothy M. Chan. Near-optimal randomized algorithms for selection in totally monotone matrices. In *SODA*, pages 1483–1495, 2021.
- 23 Timothy M. Chan and Moshe Lewenstein. Clustered integer 3sum via additive combinatorics. In *STOC*, pages 31–40, 2015.
- 24 Shucheng Chi, Ran Duan, Tianle Xie, and Tianyi Zhang. Faster min-plus product for monotone instances. In *STOC*, pages 1529–1542, 2022.
- 25 Justin Colannino, Mirela Damian, Ferran Hurtado, John Iacono, Henk Meijer, Suneeta Ramaswami, and Godfried Toussaint. An  $o(n \log n)$ -time algorithm for the restriction scaffold assignment problem. *Journal of Computational Biology*, 13(4):979–989, 2006.
- 26 Spencer Compton, Slobodan Mitrovic, and Ronitt Rubinfeld. New partitioning techniques and faster algorithms for approximate interval scheduling. *CoRR*, abs/2012.15002, 2020. URL: <https://arxiv.org/abs/2012.15002>.
- 27 Marek Cygan, Marcin Mucha, Karol Wegrzycki, and Michal Włodarczyk. On problems equivalent to  $(\min, +)$ -convolution. *ACM Trans. Algorithms*, 15(1):14:1–14:25, 2019.
- 28 Yihe Dong, Piotr Indyk, Ilya P. Razenshteyn, and Tal Wagner. Learning space partitions for nearest neighbor search. In *ICLR*, 2020.
- 29 Franziska Eberle, Nicole Megow, Lukas Nölke, Bertrand Simon, and Andreas Wiese. Fully Dynamic Algorithms for Knapsack Problems with Polylogarithmic Update Time. In *FSTTCS*, volume 213, pages 18:1–18:17, 2021.
- 30 Eduard Eiben, Daniel Lokshtanov, and Amer E Mouawad. Bisection of bounded treewidth graphs by convolutions. In *ESA 2019*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- 31 David Eppstein, Zvi Galil, and Raffaele Giancarlo. Speeding up dynamic programming. In *FOCS*, pages 488–496. IEEE Computer Society, 1988.
- 32 Guy Even, Joseph Naor, Satish Rao, and Baruch Schieber. Fast approximate graph partitioning algorithms. *SIAM J. Comput.*, 28(6):2187–2214, 1999.
- 33 Uriel Feige and Robert Krauthgamer. A polylogarithmic approximation of the minimum bisection. *SIAM J. Comput.*, 31(4):1090–1118, 2002.
- 34 Andreas Emil Feldmann and Luca Foschini. Balanced partitions of trees and applications. *Algorithmica*, 71(2):354–376, 2015.
- 35 Zvi Galil and Kunsoo Park. Dynamic programming with convexity, concavity, and sparsity. *Theor. Comput. Sci.*, 92(1):49–76, 1992.
- 36 Gramoz Goranci, Harald Räcke, Thatchaphol Saranurak, and Zihan Tan. The expander hierarchy and its applications to dynamic graph algorithms. In *SODA*, pages 2212–2228, 2021.
- 37 Chris Harrelson, Kirsten Hildrum, and Satish Rao. A polynomial-time tree decomposition to minimize congestion. In *SPAA*, pages 34–43, 2003.
- 38 Monika Henzinger, Stefan Neumann, and Andreas Wiese. Dynamic approximate maximum independent set of intervals, hypercubes and hyperrectangles. In *SoCG*, volume 164 of *LIPICs*, pages 51:1–51:14, 2020.
- 39 Dorit S Hochbaum and David B Shmoys. Using dual approximation algorithms for scheduling problems theoretical and practical results. *Journal of the ACM*, 34(1):144–162, 1987.
- 40 Yiding Hua, Rasmus Kyng, Maximilian Probst Gutenberg, and Zihang Wu. Maintaining expander decompositions via sparse cuts. *CoRR*, abs/2204.02519, 2022.
- 41 Piotr Indyk, Fabian Kuhn, and Anca Muscholl. On the fine-grained complexity of one-dimensional dynamic programming. In *ICALP 2017*, volume 80, page 21, 2017.



- 42 Klaus Jansen and Lars Rohwedder. On integer programming, discrepancy, and convolution. *arXiv preprint arXiv:1803.04744*, 2018.
- 43 Ce Jin. An improved FPTAS for 0-1 knapsack. In *ICALP*, volume 132, pages 76:1–76:14, 2019.
- 44 George Karypis and Vipin Kumar. Metis: A software package for partitioning unstructured graphs, partitioning meshes, and computing fill-reducing orderings of sparse matrices. 1997.
- 45 Donald E. Knuth. Optimum binary search trees. *Acta Informatica*, 1:14–25, 1971.
- 46 Soh Kumabe and Yuichi Yoshida. Average sensitivity of dynamic programming. In *SODA*, pages 1925–1961, 2022.
- 47 Eduardo S Laber, Wilfredo Bardales, and Ferdinando Cicalese. On lower bounds for the maximum consecutive subsums problem and the (min,+)-convolution. In *ISIT*, pages 1807–1811. IEEE, 2014.
- 48 Gary L. Miller, Richard Peng, Russell Schwartz, and Charalampos E. Tsourakakis. Approximate dynamic programming using halfspace queries and multiscale monge decomposition. In *SODA*, pages 1675–1682, 2011.
- 49 Gaspard Monge. Mémoire sur la théorie des déblais et des remblais. *Mem. Math. Phys. Acad. Royale Sci.*, pages 666–704, 1781.
- 50 Marcin Mucha, Karol Węgrzycki, and Michał Włodarczyk. A subquadratic approximation scheme for partition. In *SODA*, pages 70–88. SIAM, 2019.
- 51 Richard Peng. Approximate undirected maximum flows in  $O(m\text{polylog}(n))$  time. In *SODA*, pages 1862–1867, 2016.
- 52 Harald Räcke. Minimizing congestion in general networks. In *FOCS*, pages 43–52, 2002.
- 53 Harald Räcke and Chintan Shah. Improved guarantees for tree cut sparsifiers. In *ESA*, pages 774–785, 2014.
- 54 Harald Räcke, Chintan Shah, and Hanjo Täubig. Computing cut-based hierarchical decompositions in almost linear time. In *SODA*, pages 227–238, 2014.
- 55 Peter Sanders and Christian Schulz. Think locally, act globally: Highly balanced graph partitioning. In *SEA*, volume 7933, pages 164–175, 2013.
- 56 Hiroshi Tamura, Masakazu Sengoku, Shoji Shinoda, and Takeo Abe. Location problems on undirected flow networks. *IEICE TRANSACTIONS (1976-1990)*, 73(12):1989–1993, 1990.
- 57 Hiroshi Tamura, Masakazu Sengoku, Shoji Shinoda, and Takeo Abe. Some covering problems in location theory on flow networks. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, 75(6):678–684, 1992.
- 58 Godfried Toussaint. The geometry of musical rhythm. In *Japanese Conference on Discrete and Computational Geometry*, pages 198–212. Springer, 2004.
- 59 Godfried T Toussaint et al. A comparison of rhythmic similarity measures. In *ISMIR*, 2004.
- 60 Nithin Varma and Yuichi Yoshida. Average sensitivity of graph algorithms. In *SODA*, pages 684–703, 2021.
- 61 R. Ryan Williams. Faster all-pairs shortest paths via circuit complexity. *SIAM J. Comput.*, 47(5):1965–1985, 2018.
- 62 Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *Proceedings of the ICM*, volume 3, pages 3431–3472. World Scientific, 2018.
- 63 Andrew Chi-Chih Yao. Probabilistic Computations: Toward a Unified Measure of Complexity. In *FOCS*, pages 222–227, 1977.
- 64 F. Frances Yao. Efficient dynamic programming using quadrangle inequalities. In *STOC*, pages 429–435, 1980.

## A Organization of the Appendix

Our appendix is organized as follows:

- In Appendix B we discuss more related work.
- Appendix C introduces preliminaries.
- Appendix D presents our results for  $k$ -balanced partitioning.
- Appendix E presents our results for simultaneous source location.
- Appendix F presents our recourse lower bounds for algorithms which only maintain few solutions.
- Appendix G presents our generalization to functions with non-monotonicities and our results for  $\ell_\infty$ -necklace.
- Appendix H presents missing proofs.

## B Further Related Work

Speeding up DP algorithms is a well-studied topic, which has received attention for several decades [2, 12, 19, 22, 31, 35, 45, 48, 49, 64]. This line of work has led to several conditions, which, if satisfied, imply that the underlying DP can be solved more efficiently. These conditions include, for example, the Monge property, total monotonicity, certain convexity and concavity properties, or the Knuth–Yao quadrangle-inequality, which are often related to each other. For example, it is known that DP tables which satisfy the Monge property are also totally monotone. One of the most popular methods in this area is the SMAWK algorithm [2] which runs in near-linear time in the number of columns of the DP table if the DP table is *totally monotone*. More concretely, a DP table is totally monotone if for each submatrix  $A$  of the DP table and for every pair of consecutive rows  $i$  and  $i + 1$  in  $A$ , the minimum entry for row  $i + 1$  appears in a column that is equal to or greater than the minimum entry for row  $i$ .

However, these conditions are quite different from our conditions in Definition 8 and they are essentially incomparable. For the purpose of illustration, we will briefly argue this for total monotonicity and Definition 8; similar arguments can also be made for the Monge property and other criteria. On one hand, the totally monotone matrices do not imply that the rows of the DP table are monotone. Indeed, when the rows are monotone then finding the columns with the minimum entries is trivial (they are always in the first or last column, depending on whether we consider monotonically increasing or decreasing rows, respectively). Hence, total monotonicity does not imply our condition from Definition 8. On the other hand, the ordering of the rows is highly important for the conditions above: just swapping two rows of a totally monotone DP table can break total monotonicity. In our case, the rows can be ordered arbitrarily in the DP table, as long as their dependency graph has good properties. Hence, our property does not imply total monotonicity. This shows that these definitions are incomparable.

Recently, Varma and Yoshida [60] and Kumabe and Yoshida [46] studied the sensitivity of graph algorithms and of DP algorithms. They studied how much the solutions of such algorithms change when a random element from the input is deleted. For several problems including knapsack they showed that these algorithms have small sensitivity. However, we show in Section F that when insertions are allowed, dynamic algorithms must have high recourse or they have to maintain many different solutions.

The  $k$ -balanced graph partitioning problem has received a lot of attention in the theory community [5, 32–34]. The problem is also highly relevant in practice [18, 28, 44, 55], where algorithms for balanced graph partitioning are often used as a preprocessing step for large scale data analytics. For the special case of  $k = 2$ , this corresponds to the minimum

bisection problem and Feige and Krauthgamer [33] presented polynomial-time algorithms with polylogarithmic approximation ratios. For  $k \geq 3$ , Andreev and Räcke [5] showed that no polynomial-time algorithm can achieve a finite approximation ratio unless  $P = NP$ . They also showed how to compute a bicriteria  $(O(\log^{1.5}(n)/\epsilon^2), 1 + \epsilon)$ -approximate solution in polynomial time. Feldmann and Foschini [34] obtained a polynomial-time bicriteria  $(O(\log^{1.5}(n) \log \log n), 1 + \epsilon)$ -approximation algorithm which has the advantage that the approximation ratio does not depend on the parameter  $\epsilon$  of the partition sizes. Even et al. [32] showed that one can compute a bicriteria  $(O(\log n), 2)$ -approximation in polynomial time.

The simultaneous source location problem that we study is closely related to the source location problem introduced by Tamura et al. [56,57], in which a minimum number of sources must be selected to be able to satisfy any single demand in an undirected edge-capacitated graph. Arata et al. [7] showed that the problem is NP-hard and presented an exact algorithm for the variant with uniform vertex costs. In the simultaneous source location problem that was introduced by Andreev et al. [5] and that we study in this paper, *all* demands must be satisfied simultaneously. Andreev et al. provide an  $O(\log D)$ -approximation algorithm, where  $D$  is the sum of demands, and a matching hardness result for this problem in general graphs. They also present an exact polynomial-time algorithm when the input graph is a tree and show that this result can be extended to general graphs when the edge capacities can be violated by a  $O(\log^2 n \log \log n)$ -factor, where  $n$  is the number of vertices in the graph.

Chan [21] showed that one can consider the solutions for the 0-1 knapsack as monotone piecewise constant functions and used this insight to obtain faster algorithms. Recently, these results were improved by Jin [43] who showed how to compute a  $(1 + \epsilon)$ -approximation for 0-1 knapsack with  $n$  items in time  $\tilde{O}(n + \epsilon^{-9/4})$ . Bringmann and Cassis [15] derived faster exact algorithms for 0-1 knapsack using bounded monotone min-plus-convolution. Aouad and Segev [6] study the incremental knapsack problem, where the capacity constraint is increased over time and the goal is to find nested subsets of items which maximize the average profit; we note that this is different from our setting, where the goal is to obtain efficient update times, while the solutions may change arbitrarily over time.

An  $\ell_1$ -necklace alignment problem was first considered by Toussaint [58], motivated by computational music theory and rhythmic similarity [59]. Toussaint focused on a scenario where the beads lie at integer coordinates. Ardila et al. [8] studied the problem for binary strings. There also exist results for different distance measures between two sets of points on the real line in which not every points needs to be matched [25], as well as for computing the similarity of two melodies when they are represented as closed orthogonal chains on a cylinder [3]. Bremner et al. [14] showed that  $\ell_2$ -necklace alignment can be solved in time  $O(n \log n)$ , where  $n$  is the number of beads, using FFT. They also showed that  $\ell_\infty$ -necklace alignment can be solved using a constant number of  $(\min, +)$ -operations and obtained subquadratic-time algorithms for  $\ell_1$ - and  $\ell_\infty$ -necklace alignment.

A common subroutine that is employed when solving DPs is  $(\min, +)$ -convolution; note that this subroutine is also of high importance in all of our algorithms. The complexity of  $(\min, +)$  convolution has received significant attention in the literature [9–11, 14, 17, 20, 23, 24, 27, 41, 42, 47, 50]. It was shown that naive algorithm with running time  $O(n^2)$  can be improved to time  $n^2/2^{\Omega(\sqrt{\log n})}$  [14, 61] by a reduction to All Pairs Shortest Path [14] using Williams' algorithm for the latter [14]. However, so far, no  $O(n^{2-\epsilon})$ -time algorithm was found, which led to the *MinConv* hardness conjecture in fine-grained complexity theory [27, 41]. The conjecture is particularly appealing because it implies other conjectures such as the 3-SUM and the All-Pairs Shortest Paths conjectures, and dozens of lower bounds that follow from

them (see [27, 62]). There further exist many conditional lower bounds from the MinConv conjecture and several MinConv-equivalent problems are known, e.g., related to the knapsack problem or to subadditive sequences [27, 41], among others [1, 10, 27, 30, 41, 42, 47, 50]. There have also been improvements for efficiently approximating the  $(\min, +)$ -convolution in the case of large weights [17] for the exact  $(\min, +)$ -matrix product with bounded differences [16].

## C Preliminaries

We introduce some preliminaries that we will use in the rest of the paper. For the sake of better readability, we present some of the proofs in Appendix H. We write  $[m]$  to denote the set  $\{0, 1, \dots, m\}$ .

Throughout the paper, we will consider input graphs  $G = (V_G, E_G, \text{cap}_G)$  with  $n$  vertices and  $m$  edges, where  $\text{cap}_G : E_G \rightarrow W_\infty \cup \{\infty\}$  is a weight function that for an edge  $e \in E_G$  describes the capacity of the edge. To simplify notation we extend  $\text{cap}_G$  to all vertex pairs and define

$$\text{cap}_G(x, y) = \begin{cases} \text{cap}_G(\{x, y\}) & \{x, y\} \in E_G \\ 0 & \text{otherwise.} \end{cases} .$$

Additionally, for disjoint sets  $A, B \subseteq V_G$ , we set  $\text{cap}_G(A, B) := \sum_{(a,b) \in A \times B} \text{cap}_G(a, b)$  and  $\text{cap}_G(A) := \text{cap}_G(A, V \setminus A)$ . We drop the subscript  $G$  of the capacity function  $\text{cap}$  whenever the graph is clear from the context.

Let  $(V_T, E_T, r)$  be a rooted tree. For a vertex  $v \in V_T$  we use  $T_v$  to denote the subtree rooted at  $v$  and we say that the *degree* of  $v$  is its number of children. The *height*  $h$  of  $T$  is the length of the longest path from the root to a leaf.

### C.1 Racke Tree

A *Racke tree* [52] (or *tree cut sparsifier*)  $T = (V_T, E_T)$  for an undirected graph  $G = (V_G, E_G)$  is a weighted, rooted tree in which the leaf nodes correspond to vertices of  $G$ . For a vertex  $v \in V_T$ , we write  $V_v \subseteq V_G$  to denote the set of leaf vertices in  $T_v$ . Naturally, an edge  $e = (u, v)$  of  $T$  corresponds to a cut in  $G$ , namely to the cut formed by the set  $V_u \cap V_v$  in  $G$ . The capacity  $\text{cap}_T$  of the tree edge  $(u, v)$  is set to the capacity of this cut, i.e., to  $\text{cap}_G(V_u \cap V_v)$ .

For a graph  $H = (V_H, E_H)$  and two disjoint subsets  $A, B \subseteq V_H$ , we write

$$\text{mincut}_H(A, B) := \min_{S \subseteq V_H: A \subseteq S, B \subseteq \bar{S}} \text{cap}_H(S)$$

to denote the minimum capacity of a cut that separates  $A$  and  $B$ . By definition of the edge capacities in  $T$  we have  $\text{mincut}_T(A, B) \geq \text{mincut}_G(A, B)$  for any two disjoint subsets  $A, B \subseteq V_G$ . For the sake of completeness, we prove this property in Appendix H.5.

The goal of a Racke tree  $T$  is to approximate the cut-structure of  $G$ , i.e., to guarantee that for all disjoint sets of vertices  $A, B \subseteq V_G$ ,

$$\text{mincut}_G(A, B) \leq \text{mincut}_T(A, B) \leq q \cdot \text{mincut}_G(A, B) ,$$

for a small value  $q \geq 1$ . The parameter  $q$  is called the *quality* of the Racke tree.

In the static setting, Racke trees with polylogarithmic quality guarantees can be computed in nearly linear time [51, 54]. When larger running times are allowed, better qualities can be achieved [13, 37, 53].

► **Theorem 15** (Peng [51]). *Let  $G$  be a connected undirected graph with  $n$  vertices and  $m$  edges. Then there exist an algorithm that computes a Racke tree of height  $O(\log n)$  for  $G$  with quality  $O(\log^4 n)$  in time  $\tilde{O}(m)$ .*

Furthermore, there has recently been interest in maintaining Racke trees dynamically [36, 40]. Here, we will use a result by Goranci, Racke, Saranurak and Tan who showed that one can maintain Racke trees for unweighted graphs dynamically with subpolynomial update time.

► **Theorem 16** (Goranci, Racke, Saranurak and Tan [36]). *Let  $G$  be an undirected, unweighted graph with  $n$  vertices that is undergoing edge insertions and deletions. There exists a deterministic algorithm with amortized update time  $n^{o(1)}$  that maintains a Racke tree for  $G$  with quality  $n^{o(1)}$  and height  $O(\log^{1/6} n)$ .*

## C.2 Okay-Behaved DPs

We introduce a more general DP condition compared to the one in Definition 8 which, however, will not allow us to obtain results like Theorems 9 or 10. We will consider the same type of DP tables as in Section 2.

► **Definition 17.** *A DP is okay-behaved if it fulfills the sensitivity condition of well-behaved DPs: Suppose  $\beta > 1$  and for all  $i' \in \text{In}(i)$ , we obtain a  $\beta$ -approximation  $\text{ADP}(i', \cdot)$  of  $\text{DP}(i', \cdot)$  (as per Equation (1)). Then applying  $\mathcal{P}_i$  on the  $\text{ADP}(i', \cdot)$  yields a  $\beta$ -approximation of  $\text{DP}(i, \cdot)$ , i.e.,*

$$\text{DP}(i, \cdot) \leq \mathcal{P}_i(\{\text{ADP}(i', \cdot) : i' \in \text{In}(i)\}) \leq \beta \cdot \text{DP}(i, \cdot).$$

We also use routines  $\tilde{\mathcal{P}}_i$  to compute the DP rows  $\text{ADP}(i, \cdot)$ . Again, if for all  $i$  it holds that  $\tilde{\mathcal{P}}_i(\{\text{ADP}(i', \cdot) : i' \in \text{In}(i)\})$  is an  $\alpha$ -approximation of  $\mathcal{P}_i(\{\text{ADP}(i', \cdot) : i' \in \text{In}(i)\})$ , we say that  $\text{ADP}(1, \cdot), \dots, \text{ADP}(n, \cdot)$  is an  $\alpha$ -approximate DP solution.

In the dependency graph, we call a vertex without any incoming edges a *leaf*. The *level* of a vertex  $u$  is the length of the longest path from a leaf to  $u$ . Similar to the proof of Theorem 9 we can show the following approximation guarantee for the approximate solutions  $\text{ADP}(i, \cdot)$  and the exact solutions  $\text{DP}(i, \cdot)$ .

► **Lemma 18.** *Let  $i$  be a vertex of the dependency graph with level  $\ell$ . Then the entry  $\text{ADP}(i, \cdot)$  in the  $\alpha$ -approximate ADP-solution for a okay-behaved DP problem fulfills*

$$\text{DP}(i, \cdot) \leq \text{ADP}(i, \cdot) \leq \alpha^{\ell+1} \cdot \text{DP}(i, \cdot).$$

Next, suppose the dependency graph of the DP that we consider is derived from a tree as follows. Let  $T = (V_T, E_T, r)$  be a rooted tree with root  $r$  and height  $h$ . We assume that the children of a vertex are ordered from left to right. The dependency graph that we associate with  $T$  is simply a directed copy of  $T$  in which we direct each edge towards the root. More precisely, the dependency graph contains copies of all vertices in  $V_T$  and for each vertex  $v$  (except for  $r$ ) an edge to its parent  $p$ . Clearly, this set of edges induces a DAG in which the longest path has at most  $h$  edges. The following lemma summarizes the properties of approximate DP solutions when using this approach.

► **Lemma 19.** *Consider a rooted tree  $T = (V_T, E_T, r)$  with height  $h$ . Consider an okay-behaved DP and the ADP-solution  $\text{ADP}(i, \cdot)$  corresponding to the dependency graph described above. Assume that each  $\tilde{\mathcal{P}}_i$  is an  $\alpha$ -approximation of  $\mathcal{P}_i$  and can be computed in time at most  $t$ . Then  $\text{ADP}(r, \cdot)$  is an  $\alpha^{h+1}$ -approximation of  $\text{DP}(r, \cdot)$  and can be computed in time  $O(|V_T| \cdot t)$ .*

The main difference of this lemma together with the definition of okay-behaved DPs and Theorem 9 with well-behaved DPs is as follows. When applying Theorem 9, we only have to consider how many pieces our functions have and we do not have to bother about deriving running times bound for computing the operations on our functions (because the additional conditions from the well-behaved DPs imply good running time bounds). Here, we have to check less conditions for okay-behaved DPs (in particular, we do not have to bound the number of pieces or operations) but we have to provide our own running time analysis.

Later, when we consider dynamic algorithms, we will have to consider the scenario when the underlying tree  $T$  changes due to edge insertions and deletions (and therefore might become a forest). In that case, the dependency graph and the DP solutions  $\text{DP}(i, \cdot)$  and  $\text{ADP}(i, \cdot)$  change over time as well. The following lemma asserts that when a vertex  $i$  is affected by an edge insertion or deletion, we only have to recompute the solutions  $\text{DP}(j, \cdot)$  and  $\text{ADP}(j, \cdot)$  for vertices  $j$  that are reachable from  $i$  in the dependency graph and that there are at most  $h$  such vertices.

► **Lemma 20.** *Consider a rooted tree  $T = (V_T, E_T, r)$  with height  $h$  that is undergoing edge insertions and deletions. Then after each insertion or deletion, we can recompute an ADP-solution with the same guarantees as in Lemma 19 in time  $O(h \cdot t)$ , where  $t$  is the time it takes to compute the functions  $\tilde{P}_i$ .*

Lemma 6 already provided a way to compute the minimum of two monotone piecewise constant functions. When more than two functions are involved in the minimum computation, the following version gives improved guarantees.

► **Lemma 21.** *Let  $f_i : [0, t] \rightarrow W_\infty$ ,  $i \in \{1, \dots, k\}$  be piecewise constant functions that are either all monotonically increasing or all monotonically decreasing. Then  $f_{\min}(x) := \min_i \{f_i(x)\}$  can be computed in time  $O(\sum_i p_i \cdot \log(\sum_i p_i))$ , where  $p_i$  denotes the number of pieces of function  $f_i$ .*

We also note the following well-known lemma for sake of completeness.

► **Lemma 22.** *Let  $f_1, f_2 : [0, t] \rightarrow W_\infty$  and suppose that one of  $f_1$  and  $f_2$  is monotonically decreasing. Then  $f = f_1 \oplus f_2$  is monotonically decreasing.*

## D Balanced Graph Partitioning

In this section, we provide an algorithm for the  $k$ -balanced graph partitioning problem. In this problem, the input consists of a graph  $G = (V, E, \text{cap})$ , where  $\text{cap} : E \rightarrow W_\infty$  is a weight function on the edges, and an integer  $k$ . The goal is to find a partition  $V_1, \dots, V_k$  of the vertices such that  $|V_i| \leq \lceil |V|/k \rceil$  for all  $i$  and the weight of the edges which are cut by the partition is minimized. More formally, we want to minimize  $\text{cut}(V_1, \dots, V_k) := \sum_i \text{cap}(V_i)$ , where  $\text{cap}(V_i) = \sum_{\{u,v\} \in E \cap (V_i, V \setminus V_i)} \text{cap}(u, v)$ .

Since the above problem is NP-hard to approximate within any factor  $n^{1-\epsilon}$  for any  $\epsilon$  even on trees [34], we consider bicriteria approximation algorithms. Given a weighted graph  $G = (V, E, \text{cap})$ , we say that a partition  $V_1, \dots, V_k$  of  $V$  is an  $(\alpha, \beta)$ -approximate solution if  $|V_i| \leq \beta \lceil n/k \rceil$  for all  $i$  and  $\text{cut}(V_1, \dots, V_k) \leq \alpha \cdot \text{cut}(\text{OPT})$ , where  $\text{OPT} = (V_1^*, \dots, V_k^*)$  is the optimal solution with  $|V_i^*| \leq \lceil n/k \rceil$  for all  $i$ .

Our first main result in this section is summarized in the following theorem. We use the notation  $O'(\cdot)$  to suppress factors in  $\text{poly}(\log n, k, \log(1/\epsilon), \log \log(W))$ .

► **Theorem 2.** *Let  $\epsilon > 0$  and  $k \in \mathbb{N}$ . Let  $G = (V, E, \text{cap})$  be an undirected weighted graph with  $n$  vertices and  $m$  edges and edge weights in  $W_\infty$ . Then for the  $k$ -balanced partition problem we can compute:*

- *An  $(O(\log^4 n), 1+\epsilon)$ -approximation in time  $(k/\epsilon)^{O(\log(1/\epsilon)/\epsilon)} \cdot O'(m \cdot \log^2(W)) + (k/\epsilon)^{O(1/\epsilon^2)}$ .*
- *A  $(1+\epsilon, 1+\epsilon)$ -approximation in time  $(k/\epsilon)^{O(\log(1/\epsilon)/\epsilon)} \cdot O'(n \cdot h^2 \cdot \log^2(W)) + (k/\epsilon)^{O(1/\epsilon^2)}$  if  $G$  is a tree of height  $h$ .*
- *A  $(1, 1+\epsilon)$ -approximation in time  $(k/\epsilon)^{O(\log(1/\epsilon)/\epsilon)} \cdot O'(n^4 \cdot \log^2(W)) + (k/\epsilon)^{O(1/\epsilon^2)}$  if  $G$  is a tree.*

Furthermore, we can also extend our results to the dynamic setting in which the graph  $G$  is undergoing edge insertions and deletions. Our second main result in this section is summarized in the following theorem.

► **Theorem 3.** *Let  $\epsilon > 0$  and  $k \in \mathbb{N}$ . Let  $G = (V, E, \text{cap})$  be an undirected weighted graph with  $n$  vertices that is undergoing edge insertions and deletions. Then for the  $k$ -balanced partition problem we can maintain:*

- *An  $(n^{o(1)}, 1+\epsilon)$ -approximate solution with amortized update time  $(k/\epsilon)^{O(\log(1/\epsilon)/\epsilon)} \cdot n^{o(1)} \cdot O'(\log^2(W))$  and query time  $(k/\epsilon)^{O(1/\epsilon^2)}$  if  $G$  is unweighted.*
- *A  $(1+\epsilon, 1+\epsilon)$ -approximate solution with worst-case update time  $(k/\epsilon)^{O(\log(1/\epsilon)/\epsilon)} \cdot O'(h^3 \cdot \log^2(W))$  and query time  $(k/\epsilon)^{O(1/\epsilon^2)}$  if  $G$  is a tree of height  $h$ .*

Our DP approach is inspired by the DP of Feldmann and Foschini [34]. However, the DP cells in the algorithm of Feldmann and Foschini are not monotone and, therefore, their DP cannot directly be sped up by the fast convolution of monotone functions approach. Hence, we first simplify and generalize their DP to make it monotone such that we can apply the fast convolution of monotone functions approach.

We note that in our static and dynamic algorithms, we can output the corresponding solutions similarly to what we described after Proposition 12 for knapsack.

To obtain these results, we will first describe an exact DP in Section D.1 for the special case of binary trees. Then we will show how to compute the DP more efficiently by introducing approximation in Section D.2. In Section D.3 we show how to return a solution based on our DP table. Sections D.4 and D.5 provide extensions from binary trees to more general graphs and to the dynamic setting, respectively.

## D.1 The Exact DP

When describing the DP, we will make two assumptions. First, we assume that the input graph  $T = (V, E)$  is a *binary tree* (we show in Section D.4 how to remove this assumption). Second, we consider a slight generalization of the  $k$ -balanced partition problem on trees; we note that we did not mention this generalization in Section 4. In this generalization, we suppose that each vertex is assigned a weight by a weight function  $w: V \rightarrow \{0, 1\}$ .<sup>10</sup> For convenience we set  $w(U) = \sum_{u \in U} w(u)$  for all  $U \subseteq V$  and refer to  $w(U)$  as *the weight of the vertices in  $U$* . Now our goal will be to find a partition  $V_1, \dots, V_k$  of  $V$  such that  $w(V_i) \leq (1+\epsilon) \lceil w(V)/k \rceil$  for all  $i$  and we will compare against  $\text{OPT} = (V_1^*, \dots, V_k^*)$ , where  $\text{OPT}$  is the optimal solution with  $w(V_i^*) \leq \lceil w(V)/k \rceil$  for all  $i$ . Note that by setting  $w(v) = 1$

<sup>9</sup> We use the notation  $O'(\cdot)$  to suppress factors in  $\text{poly}(\log n, k, \log(1/\epsilon), \log \log(W))$ .

<sup>10</sup> We note that our proofs and algorithms also work for more general weight functions  $w: V \rightarrow \mathbb{R}_+$ . However, in that case the functions  $\text{DP}(v, g, \text{cut}, \cdot)$  that we will introduce later will become more complicated to compute and, therefore, we stick with the simpler case of vertex weights in  $\{0, 1\}$ .

for all  $v \in V$ , we obtain the standard  $k$ -balanced partition problem and, therefore, our variant is a strict generalization.

The reason for considering the above generalization is that later we want to use our algorithm to find a balanced partitioning of general graphs  $G = (V', E')$  using a Racke tree  $T = (V, E)$  (see Section C.1). However, the vertices  $V'$  of  $G$  are just a subset of the vertices  $V$  of the Racke tree  $T$  (since the vertices of  $G$  correspond to leaves in  $T$  and the internal nodes of  $T$  do not correspond to any vertices in  $G$ ). Thus, if we assigned weight  $w(v) = 1$  to all vertices in  $T$  and computed a balanced partitioning of  $T$ , this would not necessarily correspond to a balanced partitioning of  $G$ . Instead, later we will consider the weight function which assigns weight 1 to all leaves in  $T$  (corresponding to the vertices in  $G$ ) and weight 0 to all internal nodes of  $T$  (which can be ignored when deriving a partitioning of  $G$ ). Then each set  $V_i$  in  $T$  will correspond to a set  $V'_i$  in  $G$  with  $w(V_i) = |V'_i|$ . In particular, if  $w(V_i) \leq (1 + \epsilon)\lceil w(V)/k \rceil$  then we will obtain that  $|V'_i| \leq (1 + \epsilon)\lceil |V'|/k \rceil$  and, therefore, the sets  $V_1, \dots, V_k$  imply a balanced partition  $V'_1, \dots, V'_k$  of  $G$ .

**High-Level Description of the DP.** We start by giving a high-level description of the DP. The DP is computed bottom-up starting at the leaves of the tree  $G$  and then moving up. For each vertex  $v$ , we will compute a DP solution of minimum cost that encodes whether the edge to the parent  $p$  of  $v$  is cut and which edges shall be cut inside the subtree  $T_v$  that is rooted at  $v$ . Note that the removal of the cut edges in our solution will decompose the tree into disjoint connected components and exactly one of them contains  $v$ 's parent  $p$ . Additionally, we store information about the weight of the vertices that are still connected to the parent  $p$  (and, therefore, to the outside of  $T_v$ ) after the cut edges are removed. We will assume that when we compute the DP cell for a vertex  $v$ , we have access to the solutions for both of its children.

More concretely, when we have computed a solution for a subtree  $T_v$ , i.e., we know which edges incident to nodes in this subtree we are going to remove (note that the edge leading to the parent of  $v$  is incident to  $T_v$  and thus we consider it as part of this solution), we store the following information in the DP table. First, we store its cost, i.e., the total capacity of all edges that are incident to vertices in  $T_v$  and that are cut. As described above, we would also like to store the weight of the vertices that are connected to the parent of  $v$  and the sizes of connected components inside  $T_v$ . However, there are two difficulties: (1) We cannot store the weight of the vertices that are connected to the root exactly because this would result in a too large DP table. Instead, we store the cheapest solution in which vertices of *at most* some given weight are still connected to the parent of  $v$ . As we will see, this approach gives rise to *monotonically decreasing* functions and allows for a very efficient computation of the DP table. (2) We store implicitly the *size* of all connected components that are created after the cut edges are removed and that lie completely inside  $T_v$ . As before, storing these sizes exactly would result in a very large DP table and, therefore, we store them concisely using the concept of a *signature*. The signatures will help us to characterize the sizes of the components inside  $T_v$  very efficiently.

**Signatures.** We call a connected component in  $T_v$  *large* if it contains vertices of total weight at least  $\epsilon\lceil w(V)/k \rceil$  and otherwise we call it *small*. Let  $t = \lceil \log_{1+\epsilon}(1/\epsilon) \rceil + 1$ , and let  $M = \lceil k/\epsilon \rceil + 1$ . A *signature* is a vector  $g = (g_0, \dots, g_{t-1}) \in [M-1]^t$ . Observe that each  $g_i$  is an integer between 0 and  $M-1$  and hence there are  $M^t = (k/\epsilon)^{O(\epsilon^{-1} \log(1/\epsilon))}$  different signatures. Intuitively, an entry  $g_i$  in  $g$  tells us roughly how many components of weight  $(1 + \epsilon)^i \cdot \epsilon\lceil w(V)/k \rceil$  there are in the DP solutions that we consider. The precise definition is as follows.

Let  $\mathcal{S} = \{S_1, \dots, S_r\}$  be a set of connected components inside  $T_v$  (e.g., think of  $\mathcal{S}$  as the



components that are created after removing the cut edges in the DP solution for vertex  $v$ ). We say that a signature vector  $g = (g_0, \dots, g_{t-1}) \in [M-1]^t$  is *consistent* for  $\mathcal{S}$  if we can *match* the connected components in  $\mathcal{S}$  to entries in  $g$  as follows. For each large component  $S_j$  we let  $\ell(S_j) = \arg \min\{i \in [t] : w(S_j) \leq (1+\epsilon)^i \cdot \epsilon \lceil w(V)/k \rceil\}$ , i.e.,  $\ell(S_j)$  is the smallest number  $i$  such that  $S_j$  has weight at most  $(1+\epsilon)^i \cdot \epsilon \lceil w(V)/k \rceil$ . Let  $s_i \in [M-1]$  denote the number of times the value  $i \in [t]$  has been chosen in this process, i.e.,  $s_i = |\{j : \ell(S_j) = i\}|$ , and let  $s = (s_0, \dots, s_{t-1})$  denote the resulting vector. We say that  $g$  is *consistent* with the set of components  $\mathcal{S}$  if  $g = s$ . Thus, the above matching process can be viewed as rounding up the component sizes and counting the number of components of each size.

For  $x \in \mathbb{N}$ , we let  $e(x) \in [M-1]^t$  denote the signature of a single component with total weight  $x$ . More precisely, we set  $e(x)$  to the vector that has  $e(x)_j = 1$  for  $j = \arg \min\{j \in \mathbb{N} : x \leq (1+\epsilon)^j \cdot \epsilon \lceil w(V)/k \rceil\}$  and  $e(x)_j = 0$ , otherwise. If  $x < \epsilon \lceil w(V)/k \rceil$ , we define  $e(x) = \vec{0}$ .

### D.1.1 DP Definition

Now we describe the DP formally. An entry  $\text{DP}(v, g, \text{cut}, x) \in W_\infty$  in the DP table for a vertex  $v$  is indexed by a signature  $g$ , a Boolean value  $\text{cut}$  and  $x \in [n]$ . We will consider the tuples  $(v, g, \text{cut})$  as the rows  $\mathcal{I}$  of the DP table and  $x$  as the columns; we associate each such row with a function  $\text{DP}(v, g, \text{cut}, \cdot) : [n] \rightarrow W_\infty$ . Note that our DP has  $|V| \cdot M^t \cdot 2 = (k/\epsilon)^{O(\epsilon^{-1} \log(1/\epsilon))} \cdot n$  rows. Also, note that it has columns  $n$ ; later, even though  $x$  only takes discrete values, we will allow  $x$  to take values in  $[0, \infty)$ .

It describes the optimum cost of cutting edges incident on the subtree  $T_v$  (including the cost of maybe cutting the edge to the parent of  $v$ ). We will refer to the set of vertices in  $T_v$  that are still connected to the parent of  $v$  after the cut edges are removed as the *root component*. We impose the following conditions on  $\text{DP}(v, g, \text{cut}, x)$ :

- Once the cut edges are removed, the root component  $U \subseteq T_v$  has total weight *at most*  $x$ , i.e.,  $w(U) \leq x$ .
- If  $\text{cut}$  is set to true then the edge between  $v$  and its parent is cut, otherwise it is kept.
- The vertices inside  $T_v$  that (once the cut edges are removed) are *not* connected to the parent of  $v$  form connected components that are consistent with the signature  $g$ .

We observe that if we fix a vertex  $v$ , a signature  $g$  and a value for  $\text{cut}$ , then the resulting function  $\text{DP}(v, g, \text{cut}, \cdot)$  is monotonically decreasing in  $x$ . This will be the crucial property for the rest of the section.

► **Observation 23.** *Let  $v \in V$ ,  $g \in [M-1]^t$  be a signature and  $\text{cut} \in \{\text{true}, \text{false}\}$ . Then the function  $\text{DP}(v, g, \text{cut}, \cdot) : [0, \infty) \rightarrow \mathbb{R}_+$  is monotonically decreasing.*

**Proof.** By definition,  $\text{DP}(v, g, \text{cut}, x)$  stores the cost of the optimum solution in which the vertices in the root component have weight *at most*  $x$ . Now observe that for  $x \leq x'$ , the solution  $\text{DP}(v, g, \text{cut}, x)$  is also a feasible solution for  $\text{DP}(v, g, \text{cut}, x')$ . Therefore,  $\text{DP}(v, g, \text{cut}, \cdot)$  must be monotonically decreasing. ◀

Since the DP cells are monotonically decreasing in  $x$ , we will use the shorthand notation  $\text{DP}(v, g, \text{cut}, \infty)$  to denote the solution  $\min_x \text{DP}(v, g, \text{cut}, x)$ . Note that this minimum is obtained for the largest  $x$ -value at which  $\text{DP}(v, g, \text{cut}, \cdot)$  changes.

### D.1.2 Computing the DP

In the following, we describe how to compute  $\text{DP}(v, \cdot, \cdot, \cdot)$  exactly. For computing  $\text{DP}(v, \cdot, \cdot, \cdot)$  we simply iterate over all possible choices of  $x$ ,  $g$  and  $\text{cut}$ . Note that since each vertex has

weight in  $\{0, 1\}$ , the function  $\text{DP}(v, g, \text{cut}, \cdot)$  only changes for  $x \in [n + 1]$  (i.e., when  $x$  is an integer). Thus, we only need to consider  $n + 1$  choices for  $x$ . We conclude that to compute  $\text{DP}(v, \cdot, \cdot, \cdot)$  for a fixed vertex  $v$ , there are  $O(M^t \cdot n)$  parameter choices that we need to iterate over.

In our descriptions we use  $p$  to denote the parent of  $v$ , and  $v_l$  and  $v_r$  to denote  $v$ 's left and right child, respectively, if these exist.

**Case 1:  $v$  is a leaf.** If we cut the edge to the parent of  $v$ , then the cost is  $\text{cap}(v, p)$ , there are no vertices in the root component and  $v$  forms its own connected component with signature  $e(w(v))$ . Thus, we set  $\text{DP}(v, e(w(v)), \text{true}, x) = \text{cap}(v, p)$  for all  $x \in [0, \infty)$  and we set  $\text{DP}(v, g, \text{true}, x) = \infty$  for all  $x \in [0, \infty)$  and for all signatures  $g \neq e(w(v))$ .

Now suppose we do not cut the edge  $(v, p)$  to the parent of  $v$ . Then we do not have to pay any cost since we are not cutting any edge, the weight of vertices in the root component is  $w(v)$  and the signature is  $g = 0$  since there are no connected components in  $T_v$  that are not connected to  $p$ . Therefore, for all  $x \in [0, w(v))$  we set  $\text{DP}(v, 0, \text{false}, x) = \infty$  and for all  $x \in [w(v), \infty)$  we set  $\text{DP}(v, 0, \text{false}, x) = 0$ . For all signatures  $g \neq 0$  and all  $x \in [0, \infty)$ , we set  $\text{DP}(v, g, \text{false}, x) = \infty$ .

**Case 2:  $v$  is not a leaf.** If  $v$  is not a leaf then we assume that it has exactly two children  $v_l$  and  $v_r$  (if it has only one child, we can add a second child  $v'$  with  $w(v') = 0$ ,  $\text{cap}(v, v') = 0$  and then  $v'$  has no impact on the solution). We assume that for both  $v_l$  and  $v_r$ , we have already computed the solutions  $\text{DP}(v_l, g, \text{cut}, x)$  and  $\text{DP}(v_r, g, \text{cut}, x)$  for all possible values of  $x$ ,  $g$  and  $\text{cut}$ .

Let  $e_l = (v, v_l)$  and  $e_r = (v, v_r)$  denote the edges to the respective child and let  $e_p = (p, v)$  denote the edge to the parent  $p$  of  $v$ . In the following we distinguish four cases (A, B, C, D) depending on which of these edges we decide to cut. For each case, we compute  $\text{DP}_{\text{case}}(v, g, \text{cut}, x)$ -values,  $\text{case} \in \{A, B, C, D\}$ , which are the optimum values under the condition that we cut  $e_l$  and  $e_r$  according to the case. The final entry  $\text{DP}(v, g, \text{cut}, x)$  is then obtained by minimizing over all cases, i.e., by setting

$$\text{DP}(v, g, \text{cut}, x) = \min_{\text{case} \in \{A, B, C, D\}} \text{DP}_{\text{case}}(v, g, \text{cut}, x)$$

for all  $x, g, \text{cut}$ .

**Case A: cut  $e_l$  and  $e_r$ .** Suppose we cut  $e_l$  and  $e_r$ . Then, given  $x$  and  $g$ , we have to select subsolutions for the left and right sub-tree such that the weight of vertices that can reach  $p$  is at most  $x$  and the connected components inside are consistent with  $g$ .

First, assume we cut the edge  $e_p$ . Then the cost for cutting this edge is  $\text{cap}(v, p)$ . Furthermore, the weight of vertices inside  $T_v$  that can reach  $p$  is zero and, hence, the value of  $x$  is irrelevant by the monotonicity of  $\text{DP}(v, g, \text{cut}, \cdot)$ . Next, if we have a solution with signatures  $g_l$  and  $g_r$  in the left and right subtree, respectively, we can combine these solutions as long as  $g_l + g_r + e(w(v)) = g$  (as the vertex  $v$  forms a single component of weight  $w(v)$  since we cut both edges  $e_l$  and  $e_r$ ). Note that in the subsolution for the child  $v_l$ , the value of  $x$  does not play a role for the feasibility of the solution  $\text{DP}(v, g, \text{cut}, x)$  since the size of the root component in  $T_{v_l}$  is already encoded in  $g_l$ . Therefore, to obtain minimum cost we consider  $\text{DP}(v_l, g_l, \text{true}, \infty)$ ; by symmetry, the same holds for  $v_r$ . Therefore, we set for all  $x \in [0, \infty)$ ,

$$\text{DP}_A(v, g, \text{true}, x) = \text{cap}(v, p) + \min_{g_l + g_r = g - e(w(v))} \{\text{DP}(v_l, g_l, \text{true}, \infty) + \text{DP}(v_r, g_r, \text{true}, \infty)\}. \quad (5)$$

Second, assume we do not cut the edge to the parent  $p$ . Then there will be at least one vertex (namely  $v$ ) that can reach  $p$ . Hence,  $\text{DP}_A(v, g, \text{false}, x) = \infty$  for all signatures  $g$  and

$x \in [0, w(v))$ . For  $x \in [w(v), \infty)$ , we can combine the solutions as above and we set

$$\text{DP}_A(v, g, \text{false}, x) = \min_{g_l + g_r = g} \{ \text{DP}(v_l, g_l, \text{true}, \infty) + \text{DP}(v_r, g_r, \text{true}, \infty) \}. \quad (6)$$

**Case B: cut neither  $e_l$  nor  $e_r$ .** Next, suppose we cut neither  $e_l$  nor  $e_r$ . In this case we have to select subsolutions for  $T_{v_l}$  and  $T_{v_r}$ , where each subsolution is characterized by the upper bound  $x_l$  (resp.  $x_r$ ) and its signature  $g_l$  (resp.  $g_r$ ).

First, suppose that we cut the edge  $e_p$ . If we let  $x_l$  and  $x_r$  denote the exact weight of the root components for the subsolutions, then the vertex  $v$  will be included in a component of size  $x_l + x_r + w(v)$  afterwards. Hence, we can combine the subsolutions to a solution for signature  $g$  as long as  $g_l + g_r + e(x_l + x_r + w(v)) = g$ . Consequently we set for every  $x \in [0, \infty)$ ,

$$\text{DP}_B(v, g, \text{true}, x) = \text{cap}(v, p) + \min_{x_l, x_r, g_l + g_r = g - e(x_l + x_r + w(v))} \text{DP}(v_l, g_l, \text{false}, x_l) + \text{DP}(v_r, g_r, \text{false}, x_r).$$

Second, suppose that we do not cut  $e_p$ . Then again we have to set  $\text{DP}_B(v, g, \text{false}, x) = \infty$  for all signatures  $g$  and all  $x \in [0, w(v))$ , because the vertex  $v$  of weight  $w(v)$  can reach  $p$ . For  $x \geq w(v)$  we have to select  $x_l$  and  $x_r$  such that they sum to  $x - w(v)$  as this guarantees that vertices of weight at most  $x$  can reach the parent  $p$ . Consequently, we set for all  $x \in [w(v), \infty)$

$$\text{DP}_B(v, g, \text{false}, x) = \min_{g_l + g_r = g, x_l + x_r = x - w(v)} \text{DP}(v_l, g_l, \text{false}, x_l) + \text{DP}(v_r, g_r, \text{false}, x_r).$$

**Case C: cut  $e_l$  but not  $e_r$ .** Now suppose we cut the edge to the left child  $v_l$  but we do not cut the edge to the right child  $v_r$ . In this case,  $v$  stays connected to the root component of  $v_r$  and we need to choose a subsolution with parameters  $x_r$  and  $g_r$  for  $T_{v_r}$  and a subsolution with parameter  $g_l$  for  $T_{v_l}$ . Note that since we cut  $e_l$ , the upper bound on the weight of the root component of  $v_l$  is irrelevant as this is implicitly encoded in  $g_l$ .

First, suppose we cut  $e_p$ . If we let  $x_r$  denote the exact weight of the root component for the subsolution in  $T_{v_r}$ , then  $v$  will be included in a component of size  $x_r + w(v)$  afterwards. Hence, we can combine the subsolutions to a solution for signature  $g$  as long as  $g_l + g_r + e(x_r + w(v)) = g$ . Consequently, for every  $x \in [0, \infty)$  we set

$$\text{DP}_C(v, g, \text{true}, x) = \text{cap}(v, p) + \min_{x_r, g_l + g_r = g - e(x_r + w(v))} \text{DP}(v_l, g_l, \text{true}, \infty) + \text{DP}(v_r, g_r, \text{false}, x_r). \quad (7)$$

Second, suppose we do not cut  $e_p$ . Then we have to set  $\text{DP}_C(v, g, \text{false}, x) = \infty$  for all signatures  $g$  and all  $x \in [0, w(v))$ , because vertex  $v$  with weight  $w(v)$  can reach  $p$ . For  $x \in [w(v), \infty)$ , we have to select  $x_r \leq x - w(v)$  as this guarantees that vertices of total weight at most  $x$  can reach the parent  $p$ . Due to the monotonicity of  $\text{DP}(v_r, g_r, \text{false}, \cdot)$  we can just choose  $x_r = x - w(v)$ . Consequently, for all  $x \in [w(v), \infty)$  we set

$$\text{DP}_C(v, g, \text{false}, x) = \min_{g_l + g_r = g, x_r = x - w(v)} \text{DP}(v_l, g_l, \text{true}, \infty) + \text{DP}(v_r, g_r, \text{false}, x_r). \quad (8)$$

**Case D: cut  $e_r$  but not  $e_l$ .** Symmetric to Case C.

Next, we argue that this DP is okay-behaved, i.e., it satisfies Definition 17. In particular, we note that this DP is not well-behaved because it does not satisfy Property (4b) of

Definition 8 since in Case 2, Step B below we will have to perform too many min-operations (see Equation (11)). We will also show that the DP's dependency graph is exactly the input tree and hence the conditions of Lemma 19 are satisfied. Furthermore, all entries for a DP cell  $\text{DP}(v, \cdot, \cdot, \cdot)$  can be computed in time  $O(M^{2t}n^3)$  by simply enumerating all choices in the different min-operations above.

► **Lemma 24.** *The DP is okay-behaved and the dependency tree and the input tree  $T$  are identical. Furthermore, given a vertex  $v$ , we can compute all entries in  $\text{DP}(v, \cdot, \cdot, \cdot)$  in time  $O(M^{2t}n^3)$ .*

**Proof.** First, note that in the DP each cell  $\text{DP}(v, \cdot, \cdot, \cdot)$  only depends on the solutions of its two children. Note that these are exactly the edges which are present in the dependency graph and also in  $T$ . Therefore, the dependency graph and  $T$  are identical. Furthermore, when the input for a child solution is a  $\beta$ -approximation, the output of the DP will also be an  $\beta$ -approximation because we perform all computations exactly. Thus, the DP is also okay-behaved.

Second, let us consider the running time. Recall that for fixed  $x$ ,  $g$  and  $cut$ , we set  $\text{DP}(v, g, cut, x) = \min_{\text{case} \in \{A, B, C, D\}} \text{DP}_{\text{case}}(v, g, cut, x)$  and this quantity can be computed in time  $O(1)$  by a simple table lookup. Thus, we only have to consider the time it takes to compute  $\text{DP}_{\text{case}}(v, g, cut, x)$  for each case  $\in \{A, B, C, D\}$  and for fixed  $x$ ,  $g$  and  $cut$ .

For Case A, observe the min-operations can be computed by iterating over all  $M^t$  choices of  $g_l$  and setting  $g_r = g - e(w(v)) - g_l$  as long as  $g_r$  is a non-negative vector. Then the expressions inside the min-term can be computed by table lookup in constant time. Thus, the time is  $O(M^t)$ . For Case B, in case  $cut = \text{true}$  note that we can iterate over all choices of  $x_l$ ,  $x_r$  and iterate over  $g_l$  as described above. This takes time  $O(M^t n^2)$ . In the case  $cut = \text{false}$  we can again iterate over the  $g_l$  as above and we can iterate over all  $x_l \in [n + 1]$  and set  $x_r = x - w(v) - x_l$  as long as  $x_r \geq 0$ ; thus, the case can be solved in time  $O(M^t n)$ . For Cases C and D, we can iterate over all choices of  $x_r$  and then iterate over the  $g_l$  as above. This gives a total running time of  $O(M^t n)$ .

We conclude that for fixed  $x$ ,  $g$  and  $cut$ , the time to compute  $\text{DP}_{\text{case}}(v, g, cut, x)$  for all case  $\in \{A, B, C, D\}$  is  $O(M^t n^2)$ . Since there are  $O(n)$  choices of  $x$ ,  $M^t$  choices for  $g$  and two choices for  $cut$ , we conclude that the total running time to compute  $\text{DP}(v, \cdot, \cdot, \cdot)$  is  $O(M^{2t}n^3)$ . ◀

## D.2 The Approximate DP

In this section we show how to construct the approximate DP table in an efficient manner. For this we essentially perform the same computations as above, but instead of computing the exact solution  $\text{DP}(v, \cdot, \cdot, \cdot)$  by computing exact solutions to the cases  $\text{DP}_{\text{case}}(v, \cdot, \cdot, \cdot)$ , we compute an approximate solution  $\text{ADP}(v, \cdot, \cdot, \cdot)$  which will be the minimum of approximate solutions  $\text{ADP}_{\text{case}}(v, \cdot, \cdot, \cdot)$ , where case  $\in \{A, B, C, D\}$ .

However, there are a few crucial differences. First, for fixed  $v$ ,  $g$ ,  $cut$  and case  $\in \{A, B, C, D\}$ , we interpret  $\text{ADP}_{\text{case}}(v, g, cut, \cdot)$  as a piecewise constant function which is stored in an efficient list representation (as per Section 2). After we computed the solutions  $\text{ADP}_{\text{case}}(v, g, cut, \cdot)$ , we compute the function

$$\begin{aligned} \text{ADP}(v, g, cut, \cdot) := \\ \lceil \min\{\text{ADP}_A(v, g, cut, \cdot), \text{ADP}_B(v, g, cut, \cdot), \text{ADP}_C(v, g, cut, \cdot), \text{ADP}_D(v, g, cut, \cdot)\} \rceil_{1+\delta}, \end{aligned} \tag{9}$$

i.e., instead of just taking the minimum over the different cases, we also perform a rounding step to multiples of  $1 + \delta$ . This rounding step introduces an approximation error of  $\alpha = 1 + \delta$  but reduces the number of pieces within the piecewise constant function  $\text{DP}(v, g, \text{cut}, \cdot)$  to  $p := O(\log_{1+\delta}(W))$  according to Lemma 6 (for this to work we need to guarantee that the function to be rounded is monotone and therefore we will show that  $\text{ADP}_{\text{case}}(v, g, \text{cut}, \cdot)$  is monotone for each case  $\in \{A, B, C, D\}$ ). The second crucial difference is, of course, that we perform the above computations with values that already have been rounded, i.e., with entries from ADP instead of entries from DP. We note that Equation (9) is the *only* place in the approximate DP which is not exact; all other computations are done precisely (without any rounding) and, therefore, the approximate DP only loses a factor  $1 + \delta$ .

In order to guarantee a highly efficient implementation we rely on the following invariants for entries in the approximate DP:

1. For all  $v, g$ , and  $\text{cut}$ , the function  $\text{ADP}(v, g, \text{cut}, \cdot)$  is monotonically decreasing.
2. For all  $v, g$ , and  $\text{cut}$ , the function  $\text{ADP}(v, g, \text{cut}, \cdot)$  is piecewise constant with at most  $p := O(\log_{1+\delta}(W))$  pieces.

Note that the first property resembles the fact that for the exact DP,  $\text{DP}(v, g, \text{cut}, \cdot)$  is monotonically decreasing as per Observation 23. However, here we state this property as an invariant because there could exist approximations of  $\text{DP}(v, g, \text{cut}, \cdot)$  which are non-monotone and, therefore, we need to prove that each of our functions  $\text{ADP}(v, g, \text{cut}, \cdot)$  is indeed monotone. Note the second property follows immediately from the monotonicity and the rounding step in Equation (9) and thus we will not need to prove it in the following.

Similar to the description of the exact DP, we will now go through each of the cases and, given  $v$ , describe how to compute  $\text{ADP}(v, g, \text{cut}, \cdot)$  in time  $\tilde{O}(1)$  for all  $g$  and  $\text{cut}$ . The cases are exactly the same as for the exact DP and thus for the sake of brevity we do not repeat the correctness argument.

**Case 1:  $v$  is a leaf.** Then, we do the same in the exact case. We set  $\text{ADP}(v, e(w(v)), \text{true}, x) = \text{cap}(v, p)$  for all  $x \in [0, \infty)$  and we set  $\text{ADP}(v, g, \text{true}, x) = \infty$  for all  $x \in [0, \infty)$  and all signatures  $g \neq e(w(v))$ . Furthermore, we set  $\text{ADP}(v, 0, \text{false}, x) = \infty$  for all  $x \in [0, w(v))$  and  $\text{ADP}(v, 0, \text{false}, x) = 0$  for all  $x \in [w(v), \infty)$ . For all signatures  $g \neq 0$  and all  $x \in [0, w(v))$ , we set  $\text{ADP}(v, g, \text{false}, x) = \infty$ . Note that in all cases, the corresponding functions  $\text{ADP}(v, g, \text{cut}, \cdot)$  are monotonically decreasing and have  $O(1)$  pieces.

**Case 2:  $v$  is not a leaf.** We distinguish the same four cases as for the exact DP. Again, we will assume that  $v$  has exactly two children  $v_l$  and  $v_r$  and we let  $e_l = (v, v_l)$ ,  $e_r = (v, v_r)$  and  $e_p = (p, v)$ , where  $p$  is the parent of  $v$ .

**Case A: cut  $e_l$  and  $e_r$ .** First, suppose we cut  $e_l$  and  $e_r$ . Then, as in the exact DP, if we cut the edge to the parent of  $v$ , we wish to set

$$\begin{aligned} \text{ADP}_A(v, g, \text{true}, x) = \\ \text{cap}(v, p) + \min_{g_l + g_r = g - e(w(v))} \{ \text{ADP}(v_l, g_l, \text{true}, \infty) + \text{ADP}(v_r, g_r, \text{true}, \infty) \}. \end{aligned}$$

for all  $x \in [0, \infty)$ . Note that in the equation above, the quantities  $\text{cap}(v, p)$ ,  $\text{ADP}(v_l, g_l, \text{true}, \infty)$  and  $\text{ADP}(v_r, g_r, \text{true}, \infty)$  are simply numbers and can be viewed as a piecewise constant function with a single piece. Thus,  $\text{ADP}_A(v, g, \text{true}, \cdot)$  is a piecewise constant function with a single piece and, therefore, it is also monotonically decreasing. Hence, the invariants are satisfied for  $\text{ADP}_A(v, g, \text{true}, \cdot)$ . Furthermore,  $\text{ADP}_A(v, g, \text{true}, \cdot)$  can be computed via a sum and a minimum over monotonically decreasing piecewise functions via Lemma 6. Note that the minimum takes  $O(M^t)$  different values because it is computed by iterating over all

$g_l \in [M - 1]^t$  and setting  $g_r = g - e(w(v)) - g_l$  as long as all entries in  $g_r$  are non-negative. Since each function  $\text{ADP}(v_l, g_l, \text{true}, \cdot)$  has  $O(p)$  pieces according to our invariants, we can compute the value  $\text{ADP}(v_l, g_l, \text{true}, \infty)$  in time  $O(1)$ ; the same holds for  $\text{ADP}(v_r, g_r, \text{true}, \infty)$ . Thus, computing  $\text{ADP}_A(v, g, \text{true}, \cdot)$  takes time  $O(M^t)$ .

Next, suppose we do not cut the edge to the parent of  $v$ . Then, as in the exact DP, we wish to set:

$$\text{ADP}_A(v, g, \text{false}, x) = \min_{g_l + g_r = g} \{ \text{ADP}(v_l, g_l, \text{true}, \infty) + \text{ADP}(v_r, g_r, \text{true}, \infty) \}$$

for all  $x \in [0, \infty)$ . Then by the same arguments as above,  $\text{ADP}_A(v, g, \text{false}, \cdot)$  is a piecewise constant monotonically decreasing function with a single piece. It can be computed in time  $O(M^t)$  as described above.

**Case B: cut neither  $e_l$  nor  $e_r$ .** Now suppose we do not cut any edge to the children.

If we do not cut the edge to the parent of  $v$ , we proceed similar to the exact DP. We start by setting  $\text{ADP}_B(v, g, \text{false}, x) = \infty$  for all  $x \in [0, w(v))$ . Next, for  $x \in [w(v), \infty)$  we wish to set

$$\begin{aligned} \text{ADP}_B(v, g, \text{false}, x) &= \min_{g_l + g_r = g, x_l + x_r = x - w(v)} \text{ADP}(v_l, g_l, \text{false}, x_l) + \text{ADP}(v_r, g_r, \text{false}, x_r) \\ &= \min_{g_l + g_r = g} \min_{x_l + x_r = x - w(v)} \text{ADP}(v_l, g_l, \text{false}, x_l) + \text{ADP}(v_r, g_r, \text{false}, x_r). \end{aligned} \tag{10}$$

Note that for fixed  $g_l$  and  $g_r$ , the inner min-operation in the second line describes a  $(\min, +)$ -convolution due to the constraint  $x_l + x_r = x - w(v)$ . Therefore, in the inner min-operation we compute a convolution  $\text{ADP}(v_l, g_l, \text{false}, \cdot) \oplus \text{ADP}(v_r, g_r, \text{false}, \cdot)$  and shift the result by  $w(v)$  via the shift operation from Lemma 6 (where for  $x \in [0, w(v))$  we set  $\text{ADP}_B(v, g, \text{false}, x) = \infty$ ). We need time  $O(p^2 \log p)$  for computing the convolution according to Lemma 7. To compute the outer minimum in Equation (10), we iterate over all  $g_l \in [M - 1]^t$  and thus perform  $O(M^t)$  minimum computations over piecewise constant functions with at most  $p^2$  pieces. Hence, we need time  $O(M^t p^2 \log(M^t p^2))$  according to Lemma 21. By Lemma 22,  $\text{ADP}_B(v, g, \text{false}, \cdot)$  is monotonically decreasing since it is the minimum over convolutions of two monotonically decreasing functions.

If we cut the edge to the parent of  $v$ , then for all  $x \in [0, \infty)$  we would like to set

$$\begin{aligned} \text{ADP}_B(v, g, \text{true}, x) &= \\ \text{cap}(v, p) + \min_{x_l, x_r, g_l + g_r = g - e(x_l + x_r + w(v))} &\text{ADP}(v_l, g_l, \text{false}, x_l) + \text{ADP}(v_r, g_r, \text{false}, x_r). \end{aligned}$$

Note that here we need to be careful as the range of  $g_l$  and  $g_r$  depends on the choice of  $x_l + x_r$ . Since there are  $\Omega(n)$  possible values for  $x_l + x_r$ , we cannot afford to iterate over all values that  $x_l + x_r$  can take. Instead, we will show that we only need to consider  $O(\log(k/\epsilon)/\epsilon)$  different pairs  $(x_l, x_r)$  by exploiting the monotonicity of  $\text{ADP}(v_l, g_l, \text{false}, \cdot)$  and  $\text{ADP}(v_r, g_r, \text{false}, \cdot)$ .

First, observe that we can assume  $x_l \leq w(T_{v_l})$  and  $x_r \leq w(T_{v_r})$ : increasing the upper bounds on the weight of the root component further would mean that the root component contains more weight than *all* vertices inside the sub-tree, which is impossible. Thus,  $x_l + x_r + w(v) \in [1, w(V)]$ .

Second, we partition the interval  $[1, w(V)]$  into  $O(\log(k/\epsilon)/\epsilon)$  intervals. We have intervals  $I_j = (\xi_{j-1}, \xi_j]$  with  $\xi_j = (1 + \epsilon)^j \lceil w(V)/k \rceil$  for all  $j = 1, \dots, \log_{1+\epsilon}(k/\epsilon)$ . In addition, we add an ‘‘interval’’  $I_0 := [\epsilon \lceil w(V)/k \rceil, \epsilon \lceil w(V)/k \rceil]$  and the interval  $I_{-1} := [1, \epsilon \lceil w(V)/k \rceil]$ . We set  $\xi_0 = \epsilon \lceil w(V)/k \rceil$  and we set  $\xi_{-1}$  to the largest integer that is less than  $\epsilon \lceil w(V)/k \rceil$ . Observe

that for all  $j \geq -1$  and  $x \in I_j$ , we have  $e(x) = e(\xi_j)$ , i.e., the value of  $e(x)$  does not change on in the interval  $I_j$ . Below, this property will allow us to separate the conditions on  $x_l + x_r$  and on  $g_l + g_r$ .

Now we can rewrite the above expression as

$$\text{ADP}_B(v, g, \text{true}, x) = \text{cap}(v, p) + \min_j \min_{x_l + x_r + w(v) \in I_j} \min_{g_l + g_r = g - e(\xi_j)} \text{ADP}(v_l, g_l, \text{false}, x_l) + \text{ADP}(v_r, g_r, \text{false}, x_r).$$

Third, note that now the two min-operations only depend on the choice of  $j$  and, importantly, the minimum over  $g_l$  and  $g_r$  does not depend on the choice of  $x_l + x_r$  anymore. Therefore, we can swap the order of the two min-operations. Furthermore, since  $\text{ADP}_B(v, g, \text{false}, x)$  is monotonically decreasing with  $x$ , we can restrict the choice of  $x_l$  and  $x_r$  such that  $x_l + x_r + w(v)$  is the largest number in the corresponding interval  $I_j$ , i.e.,  $x_l + x_r + w(v) = \xi_j$ . Thus,

$$\begin{aligned} \text{ADP}_B(v, g, \text{true}, x) &= \text{cap}(v, p) + \\ \min_j \min_{g_l + g_r = g - e(\xi_j)} \min_{x_l + x_r + w(v) = \xi_j} &\text{ADP}(v_l, g_l, \text{false}, x_l) + \text{ADP}(v_r, g_r, \text{false}, \xi_j - x_l - w(v)). \end{aligned} \tag{11}$$

Next, we explain how the above expression can be computed efficiently. Let us first argue how we can efficiently compute the inner min-operation of the above expression. We start by observing that this min-operation is *not* a convolution since in the constraint we sum up to  $\xi_i$  which is a constant (rather than to the variable  $x$ ). Now recall that  $\text{ADP}(v_l, g_l, \text{false}, \cdot)$  and  $\text{ADP}(v_r, g_r, \text{false}, \cdot)$  are piecewise constant functions with  $O(p)$  pieces by our invariants. Since  $x_l, x_r \geq 0$  this implies that there are only  $O(p^2)$  choices for  $x_l$  and  $x_r$  such that  $x_l, x_r \in I_j$  and either a new piece starts in  $\text{ADP}(v_l, g_l, \text{false}, x_l)$  or in  $\text{ADP}(v_r, g_r, \text{false}, x_r)$ . Thus, we can iterate over all these pairs  $(x_l, x_r)$  and evaluate  $\text{ADP}(v_l, g_l, \text{false}, x_l) + \text{ADP}(v_r, g_r, \text{false}, x_r)$ , where  $x_r = \xi_j - x_l - w(v)$ . Thus, we can compute the inner min-operation in time  $O(p^2 \log p)$ .

Next, we can compute the outer two min-operations by simply iterating over  $j$  and all choices for  $g_l$  and setting  $g_r = g - e(\xi_j) - g_l$  as above in  $O(M^t \cdot \log(k/\epsilon)/\epsilon)$  iterations. Hence, we obtain a running time of  $O(M^t p^2 \log p \cdot \log(k/\epsilon)/\epsilon)$ . We note that this is the step which makes the okay-behaved rather than well-behaved (since it violates Property (4b) of Definition 8).

Finally, we note that as  $\text{ADP}_B(v, g, \text{true}, x)$  is independent of  $x$ , it is a constant. Thus,  $\text{ADP}_B(v, g, \text{true}, x)$  is a piecewise constant function with a single piece and it is monotonically decreasing.

**Case C: cut  $e_l$  but not  $e_r$ .** Now suppose we cut the edge to the left child but not to the right child.

First assume that we cut the edge to the parent of  $v$ . As in the exact DP, for all  $x \in [0, \infty)$  we want to set

$$\text{ADP}_C(v, g, \text{true}, x) = \text{cap}(v, p) + \min_{x_r, g_l + g_r = g - e(x_r + w(v))} \text{ADP}(v_l, g_l, \text{true}, \infty) + \text{ADP}(v_r, g_r, \text{false}, x_r).$$

As in the previous case, observe that in the minimum the constraint  $g_l + g_r = g - e(x_r + w(v))$  depends on the choice of  $x_r$ . Thus, we rewrite the above equation analogously to the previous

case:

$$\begin{aligned}
& \text{ADP}_C(v, g, \text{true}, x) \\
&= \text{cap}(v, p) + \min_j \min_{x_r+w(v) \in I_j} \min_{g_l+g_r=g-e(\xi_j)} \text{ADP}(v_l, g_l, \text{true}, \infty) + \text{ADP}(v_r, g_r, \text{false}, x_r) \\
&= \text{cap}(v, p) + \min_j \min_{g_l+g_r=g-e(\xi_j)} \min_{x_r+w(v) \in I_j} \text{ADP}(v_l, g_l, \text{true}, \infty) + \text{ADP}(v_r, g_r, \text{false}, x_r) \\
&= \text{cap}(v, p) + \min_j \min_{g_l+g_r=g-e(\xi_j)} \text{ADP}(v_l, g_l, \text{true}, \infty) + \text{ADP}(v_r, g_r, \text{false}, \xi_j - w(v)),
\end{aligned}$$

where in the last step we used that  $\text{ADP}(v_r, g_r, \text{false}, \cdot)$  is monotonically decreasing. The evaluation of the function values of the two piecewise constant functions with  $O(p)$  pieces can be done in time  $O(\log(p))$ . Furthermore, by exhaustively enumerating all choices for  $j$  and proceeding for  $g_l$  and  $g_r$  as above, we obtain  $O(M^t \log(k/\epsilon)/\epsilon)$  iterations giving a total running time of  $O(M^t \log p \log(k/\epsilon)/\epsilon)$ . As before,  $\text{ADP}_C(v, g, \text{true}, \cdot)$  is a constant (since the computation does not depend on  $x$ ) and therefore it has only a single piece and it is monotonically decreasing.

Next, suppose we do not cut the edge to the parent of  $v$ . Then we set  $\text{DP}_C(v, g, \text{false}, 0) = \infty$  for all signatures  $g$  and all  $x \in [0, w(v))$ . For all  $x \in [w(v), \infty)$ , we set

$$\text{ADP}_C(v, g, \text{false}, x) = \min_{g_l+g_r=g} \text{ADP}(v_l, g_l, \text{true}, \infty) + \text{ADP}(v_r, g_r, \text{false}, x - w(v)).$$

Note that inside the min-operation, the first term is a constant and the second term is a piecewise constant function that is shifted by  $w(v)$ . Furthermore, the minimum is taken over  $O(M^t)$  piecewise constant functions (one for each choice of  $g_l$  by the same argument as above). We can perform the addition and shift operation via Lemma 6 (time  $O(p \log p)$  per application). Then we perform a minimum operation over  $M^t$  functions where each function has just  $p$  pieces. This can be done in time  $O(M^t p \log(M^t p))$  by Lemma 21. In total we get a running time of  $O(M^t p \log(M^t p))$ .

**Case D: cut  $e_r$  but not  $e_l$ .** Symmetric to Case C.

We conclude this subsection with the following lemma which summarizes the properties of the approximate DP computation. The lemma follows immediately from the above discussion.

► **Lemma 25.** *The approximate DP computes a  $(1 + \delta)$ -approximate DP solution and the dependency tree and the input tree  $T$  are identical. Given a vertex  $v$ , a signature  $g$  and value  $\text{cut} \in \{\text{true}, \text{false}\}$ , we can compute the corresponding approximate DP entry  $\text{ADP}(v, g, \text{cut}, \cdot)$  in time  $O(M^t p^2 \log(M^t p) \log(k/\epsilon)/\epsilon)$ .*

**Proof.** The approximation ratio of the approximate DP is  $(1 + \delta)$ -approximate because, as we pointed out earlier, we only use exact computations except in the rounding step in Equation (9). Thus, we only use  $(1 + \delta)$ -factor in the computation.

The claim about the running time follows immediately from the discussion above the lemma, where we already analyzed the running times for all steps. ◀

### D.3 Computing the Result

In this section we describe how the previously described DPs can be used to extract the result for the  $k$ -balanced partition problem. Recall that we consider the generalized version of the  $k$ -balanced partition problem, where each vertex  $v$  has a weight  $w(v) \in \{0, 1\}$  (see Section D.1 for the definition).



We focus on the *value version* of the problem in which we only need to output an approximation of the *value* of the optimal cut OPT but we do not have to return the actual partition  $V_1, \dots, V_k$  that obtains this cut value. We note, however, that by analyzing the DP solution from top to bottom, we could also construct a concrete partition  $V_1, \dots, V_k$  in time  $\tilde{O}(n)$  that achieves the cut value which is returned by the value version.

**Feasible Signatures.** Before we describe our algorithm, we first need to introduce the notion of *feasible* signatures. More concretely, recall that in Section D.1 we introduced signatures as a succinct way of storing the sizes of connected components in a solution. Now, feasible signatures will refer to signatures in which the connected components can be partitioned such that we obtain a nearly  $k$ -balanced partitioning of the vertices. We make this intuition more formal below.

For every signature  $g = (g_0, \dots, g_{t-1}) \in [M-1]^t$ , we say that its associated machine scheduling instance<sup>11</sup>  $I(g)$  is the instance which contains exactly  $g_i$  jobs of size  $(1 + \epsilon)^i \cdot \epsilon \lceil w(V)/k \rceil$  of all  $i$ . We say that  $g$  is a *feasible* signature if the jobs in  $I(g)$  can be scheduled on  $k$  machines with makespan at most  $(1 + \epsilon) \lceil w(V)/k \rceil$ . Later, we will identify the machines of the scheduling problems with partitions in the  $k$ -balanced partitioning solution and the jobs with connected components. In this way, we will be able to ensure the balance constraints of the  $k$ -balanced partitioning solution.

**Algorithm.** We now describe our two static algorithms for binary trees. The only difference between the algorithms is whether to use the exact DP from Section D.1 or the approximate DP from Section D.2; we will refer to these algorithms as the *exact* and the *approximation* algorithm, respectively. We assume that the input is an error parameter  $\epsilon > 0$  and a rooted, weighted tree  $T = (V, E, \text{cap})$  with root  $r$  and vertex weights  $w(v) \in \{0, 1\}$  for which we wish to solve the  $k$ -balanced partitioning problem.

First, our algorithm augments  $T$  by adding a *fake root*  $r'$ . We make  $r'$  the parent of  $r$  and set  $w(r') = 0$  and  $\text{cap}(r, r') = 0$ . Then we compute the DP bottom-up as described in Section C.2, where we interpret  $T$  as its own dependency graph. In the exact algorithm, we use the DP from Section D.1, and in the approximation algorithm, we use the DP from Section D.2.

Second, we compute the set of all *nearly* feasible signatures. To obtain this set, we enumerate all  $M^t$  signatures and for each of them, we check whether it is nearly feasible or not. We do this as follows. For each signature  $g$ , we construct the machine scheduling instance  $I(g)$  and run the PTAS by Hochbaum and Shmoys [39] for this problem with approximation ratio  $1 + \bar{\epsilon}$  and running time  $(N/\bar{\epsilon})^{O(1/\bar{\epsilon}^2)}$ , where  $N$  denotes the total number of jobs in  $I(g)$  and we will see later that  $N$  is a constant if  $k$ ,  $\epsilon$  and  $\bar{\epsilon}$  are constants. We add a signature  $g$  to the set of nearly feasible signatures if the returned makespan for  $I(g)$  is at most  $(1 + \bar{\epsilon})(1 + \epsilon) \lceil w(V)/k \rceil$ . We note that by using the PTAS, the set that we compute can potentially contain some signatures which are infeasible but they still do not violate the balance constraint too much.

Third, we consider the entries in the DP table at the (true) root  $r$  of the tree for the case that the edge to its (artificial) parent is not cut (recall that we added an edge of weight 0 from the true root  $r$  to the fake root  $r'$  and so cutting it does not incur any cost), i.e., we consider the DP entries  $\text{DP}(r, \cdot, \text{true}, w(V))$  or  $\text{ADP}(r, \cdot, \text{true}, w(V))$  depending on whether we are in the approximate or in the exact case. We iterate over all *feasible* signature vectors

<sup>11</sup> Recall that in the makespan minimization problem with identical machines, the input consists of a set of  $N$  jobs of sizes  $s_1, \dots, s_N$  and an integer  $k$ . The goal is to find an assignment of the jobs to  $k$  machines such that the makespan is minimized. Here, the makespan refers to maximum load of all  $k$  machines.

$g$  and then take the minimum value that we have seen.

We conclude the algorithms' guarantees in the following proposition. We note that for constant  $k$ ,  $\epsilon$ ,  $\bar{\epsilon}$  and  $W$ , the running time of the exact algorithm is  $\tilde{O}(n^4)$  and the running time of the approximation algorithm simplifies to  $\tilde{O}(n \cdot h^2)$ , where  $h$  is the height of the input tree. Thus, for trees of height  $\tilde{O}(1)$ , the approximation algorithm is very efficient and runs in time  $\tilde{O}(n)$ .

► **Proposition 26.** *Let  $\epsilon, \bar{\epsilon} > 0$  and  $k \in \mathbb{N}$ . Let  $T = (V, E, \text{cap})$  be a rooted binary tree that has edge weights  $\text{cap}(e)$  and vertex weights  $w(v) \in \{0, 1\}$ . Then:*

- *The exact algorithm obtains a bicriteria  $(1, (1 + \bar{\epsilon})(1 + \epsilon))$ -approximation for the  $k$ -balanced partitioning problem on  $T$  in time  $O(M^{2t}n^4)$ .*
- *The approximation algorithm obtains a bicriteria  $(1 + \epsilon, (1 + \bar{\epsilon})(1 + \epsilon))$ -approximation for the  $k$ -balanced partitioning problem on  $T$  in time  $O(nh^2 \cdot M^{2t} \log^2(W) \log(k/\epsilon) \log(M^t h \log(W)/\epsilon)/\epsilon^3) + M^t(k/(\epsilon\bar{\epsilon}))^{O(1/\epsilon^2)}$ , where  $h$  denotes the height of  $T$ .*

**Proof.** To prove the proposition, we need to argue about the approximation ratios of the algorithms and we also need to prove that the partitioning does not violate the balance constraints. We will also need to analyze the running times.

We start by analyzing the balance constraints. We show that in the solution returned by the algorithm, the connected components  $V_1, \dots, V_k$  can be partitioned such that  $w(V_i) \leq (1 + \bar{\epsilon})(1 + \epsilon)\lceil w(V)/k \rceil$  for all  $i = 1, \dots, k$ .

Consider the DP entry  $\text{DP}(r, g, \text{true}, w(V))$  for the (true) root  $r$ , where the edge to the parent is cut and any signature vector  $g$  that is in the set of nearly feasible signatures that we computed. Then this corresponds to the cost of some partition of  $T = T_r$  where, after removing the cut edges, the large connected components  $\mathcal{S}$  in  $T_r$  can be matched to entries in  $g$  such that:

- a component  $S \in \mathcal{S}$  is matched to entry  $g_i$  with  $|S| \leq (1 + \epsilon)^i \epsilon \lceil w(V)/k \rceil$  and
- exactly  $g_i$  components are matched to  $g_i$ .

Hence, we can obtain a partitioning  $V_1, \dots, V_k$  as follows. First, we compute the  $(1 + \bar{\epsilon})$ -approximate solution of  $I(g)$  in which (by assumption on  $g$ ) the makespan is at most  $(1 + \bar{\epsilon})(1 + \epsilon)\lceil w(V)/k \rceil$ . This gives us an assignment of jobs to machines. Now we identify components with jobs and the sets  $V_i$  with machines and obtain an assignment of the large components to the  $V_i$ . In particular, each  $V_i$  receives large components for which the (rounded) weights sum to at most  $(1 + \bar{\epsilon})(1 + \epsilon)\lceil w(V)/k \rceil$ . Now we need to assign the small components in the algorithm's solution. These can be assigned greedily by always assigning a small component (of weight less than  $\epsilon \lceil n/k \rceil$ ) to set  $V_i$  of (currently) smallest weight. In the end, all  $V_i$  will have weight at most  $(1 + \bar{\epsilon})(1 + \epsilon)\lceil w(V)/k \rceil$  (this follows from the standard argument that, when considering exact component weights, on average each server has makespan at most  $w(V)/k$  and thus there will always be a server of makespan at most  $w(V)/k$  to which the current small component can be assigned without violating the capacity constraint). This means if the algorithm returns an objective function value then there is a partition  $V_1, \dots, V_k$  with the same objective function value that is nearly feasible, i.e., that satisfies  $w(V_i) \leq (1 + \bar{\epsilon})(1 + \epsilon)\lceil w(V)/k \rceil$  for all  $i = 1, \dots, k$ .

Next, let us consider the approximation ratios of the algorithms. Consider the optimum partition  $\text{OPT} = (V_1^*, \dots, V_k^*)$  that minimizes  $\text{cut}(V_1^*, \dots, V_k^*)$  such that  $w(V_i^*) \leq \lceil w(V)/k \rceil$  for all  $i$ . We first argue that  $\text{OPT}$  gives rise to a DP entry with a feasible signature and cost  $\text{OPT}$  in the exact DP. To see this, take the optimum partition  $V_1^*, \dots, V_k^*$  and round up the weight of every large connected component to the next value of the form  $(1 + \epsilon)^i \cdot \epsilon \lceil n/k \rceil$ . Let  $g = (g_0, \dots, g_{t-1}) \in [M - 1]^t$  be the signature where  $g_i$  denotes the

number of large components in OPT whose rounded weight is  $(1 + \epsilon)^i \cdot \epsilon \lceil n/k \rceil$ . Note that since  $w(V_i^*) \leq \lceil w(V)/k \rceil$  for all  $i$ , the total *rounded* weight of components in  $V_i^*$  is at most  $(1 + \epsilon) \cdot \lceil w(V)/k \rceil$  as component weights are increased at most by a  $(1 + \epsilon)$ -factor. Hence, the constructed signature vector  $g$  is feasible because the partition  $V_1^*, \dots, V_k^*$  gives rise to a feasible solution for  $I(g)$ . Furthermore, the rounding did not have any effect on the objective function value and, thus, OPT gives rise to a DP entry with a feasible signature and cost OPT in the exact DP. This implies that the optimum value for the exact DP is at most  $cut(V_1^*, \dots, V_k^*)$ . Together with the above claim that the DPs approximately satisfy the balance constraint, we obtain that the exact algorithm computes a bicriteria  $(1, (1 + \bar{\epsilon})(1 + \epsilon))$ -approximation.

Now let us turn to the approximation ratio of the approximation algorithm from Section D.2. Recall that by Lemma 24 the exact DP is okay-behaved and in Lemma 25 we show that in each step the approximation algorithm loses a factor of at most  $1 + \delta$  at every level of the tree  $T$ . Now, we can apply Lemma 19 to obtain that the approximation in the root is  $(1 + \delta)^{h+1}$ , where  $h$  is the height of the tree  $T$ . Thus, the approximation ratio of the approximate DP is  $1 + \epsilon$  if we set  $\delta = \ln(1 + \epsilon)/(h + 1)$  since then  $(1 + \delta)^{h+1} \leq \exp(\delta(h + 1)) = 1 + \epsilon$ . Since the notion of approximation from Lemma 19 holds for all functions of the form  $ADP(r, g, \text{true}, \cdot)$  and all possible values of  $x$ , we obtain that the approximation algorithm computes a bicriteria  $(1 + \epsilon, (1 + \bar{\epsilon})(1 + \epsilon))$ -approximation.

We conclude the proof of the proposition by considering the running times of the algorithms. Note that w.r.t. running time, both algorithms only differ by how long it takes to fill the DP cells and the time for computing the solution is the same.

Let us first consider the time for computing the solution as per Section D.3. First, let us consider the time for solving the PTAS which is  $(N/\bar{\epsilon})^{O(1/\bar{\epsilon}^2)}$ , where  $N$  denotes the total number of jobs. Note that in our case there are at most  $N \leq k(1 + 1/\epsilon)$  jobs: each job has size at least  $\epsilon \lceil n/k \rceil$  and therefore a machine can take at most  $1 + 1/\epsilon$  jobs in an optimum solution. Hence, if we have more than  $k(1 + 1/\epsilon)$  jobs, a PTAS can directly reject the instance and declare it infeasible. Thus, the time for running the PTAS a single time is  $(k/(\epsilon\bar{\epsilon}))^{O(1/\bar{\epsilon}^2)}$ . Since we have to run the PTAS for each of the  $M^t$  signatures, the total time for finding the nearly feasible configurations is  $M^t (k/(\epsilon\bar{\epsilon}))^{O(1/\bar{\epsilon}^2)}$ .

Finally, let us consider the time for filling the DP cells. For the exact DP, Lemma 24 states that filling a cell  $DP(v, \cdot, \cdot, \cdot)$  takes time  $O(M^{2t} n^3)$ . Then, by applying Lemma 19, the total time to compute all DP cells is  $O(M^{2t} n^4)$ . For the approximate DP, it takes time  $O(M^t p^2 \log(M^t p) \log(k/\epsilon)/\epsilon)$  to fill a single DP cell  $ADP(v, g, cut, \cdot)$  by Lemma 25. Since there are  $M^t$  choices for  $g$  and by again applying Lemma 19, we obtain that the total running time for filling the approximate DP table is  $O(nM^{2t} p^2 \log(M^t p) \log(k/\epsilon)/\epsilon)$ . Since in Section D.2 we picked the number of pieces to be  $p = O(\log_{1+\delta}(W))$  and above we picked  $\delta = O(\epsilon/h)$ , the running time is upper bounded by  $O(nM^{2t} \cdot (1/\epsilon \cdot h \log W)^2 \cdot \log(k/\epsilon)/\epsilon \cdot \log(M^t h \log(W)/\epsilon)) = O(nh^2 \cdot M^{2t} \log^2(W) \log(k/\epsilon) \log(M^t h \log(W)/\epsilon)/\epsilon^3)$ . ◀

## D.4 Extension to General Graphs

Now we generalize the results of Proposition 26 from binary trees to general graphs.

We start with the generalization to general graphs in which we will make use of Räcke trees (see Section C.1). Since Räcke trees might be non-binary, we now introduce the notion of *binarized Räcke trees* which essentially describe a way of turning a non-binary Räcke tree into a binary tree that is very similar to a Räcke tree. Later, the binarized Räcke trees will allow us to apply Proposition 26 on them.

► **Definition 27** (Binarized Racke Tree). *Let  $G = (V_G, E_G, \text{cap}_G)$  be a weighted graph. We say that a weighted, rooted tree  $T = (V_T, E_T, \text{cap}_T)$  is a binarized Racke tree for  $G$  if the following properties hold:*

- $T$  is a rooted binary tree.
- $V_G \subseteq V_T$ .
- All edges in  $T$  have weights in  $W_\infty$ .
- Let  $T'$  be the tree that is obtained by contracting all edges with weight  $\infty$  in  $T$ . Then  $T'$  is a Racke tree for  $G$ .

We call the tree  $T'$  from the last bullet point the corresponding (non-binarized) Racke tree of  $T$ . We say that  $T$  has quality  $q$  if the corresponding Racke tree  $T'$  has quality  $q$ .

Next, we observe that each cut in  $T$  of finite cost corresponds to a cut in the corresponding Racke tree  $T'$  and vice versa. Therefore, cuts of finite cost in  $T'$  approximate the cut structure of the initial graph  $G$ . We make this more formal in following observation.

► **Observation 28.** *Let  $G = (V_G, E_G, \text{cap}_G)$  be a weighted graph and let  $T = (V_T, E_T, \text{cap}_T)$  be a binarized Racke tree for  $G$  with quality  $q$ . Then for all disjoint subsets  $A, B \subseteq V_G$  it holds that  $\text{mincut}_G(A, B) \leq \text{mincut}_T(A, B) \leq q \cdot \text{mincut}_G(A, B)$ .*

**Proof.** Let  $T'$  be the corresponding (non-binarized) Racke tree of  $T$  and consider two disjoint subsets of vertices  $A, B \subseteq V_G$ . We show that  $\text{min}_T(A, B) = \text{mincut}_{T'}(A, B)$ . Then the observation follows immediately since  $T$  has quality  $q$  (by assumption) and, therefore,  $T'$  is a Racke tree for  $G$  with quality  $q$  which satisfies the property from the observation.

Since  $T'$  can be obtained from  $T$  only by contracting edges, we have  $\text{mincut}_{T'}(A, B) \geq \text{mincut}_T(A, B)$ . Next, let us argue that  $\text{mincut}_T(A, B) \geq \text{mincut}_{T'}(A, B)$ . First, note that  $\text{mincut}_{T'}(A, B) \leq q \cdot \text{mincut}_G(A, B) < \infty$ . Since we contract only edges with weight  $\infty$  to go from  $T$  to  $T'$ ,  $T$  does not contain any cut with finite cost that is not contained in  $T'$ . Therefore,  $\text{mincut}_T(A, B) \geq \text{mincut}_{T'}(A, B)$ . ◀

Additionally, we show that we can compute a binarized Racke tree of good quality in nearly-linear time.

► **Lemma 29.** *Let  $G = (V_G, E_G, \text{cap}_G)$  be a weighted graph with  $n$  vertices and  $m$  edges. We can compute a binarized Racke tree  $T = (V_T, E_T, \text{cap}_T)$  with  $O(n)$  vertices, height  $O(\log^2 n)$  and quality  $O(\log^4 n)$  in time  $\tilde{O}(m)$ .*

**Proof.** Let  $T' = (V_{T'}, E_{T'}, \text{cap}_{T'})$  be the Racke tree for  $G$  from Theorem 15 that can be computed in time  $\tilde{O}(m)$ . First, note that  $T'$  has  $n_{T'} := O(n)$  vertices and height  $O(\log n)$ . Second, note that  $T'$  can have unbounded degree. Therefore, we will show how to compute a binarized Racke tree  $T$  that has  $T'$  as its corresponding (non-binarized) Racke tree. We do so replacing in  $T'$  each vertex  $u$  by a balanced binary tree  $\tau_u$  with  $\text{deg}(u)$  leaves, where  $\text{deg}(u)$  denotes the number of children of  $u$ . The internal edges of  $\tau_u$  will have weight  $\infty$  and the edges connecting subtrees  $\tau_u$  and  $\tau_v$ ,  $u \neq v$ , in  $T$  will correspond to the edges in  $T'$  and will have the same (finite) weight as in  $T'$ . We will see that by contracting all edges with weight  $\infty$  in  $T$ , we will obtain  $T'$ . We now elaborate on this process.

We construct  $T$  as follows. First, we compute  $T'$  as per the algorithm from Theorem 15. Now we construct  $T$  as follows. For each vertex  $u \in V_{T'}$ , we add a balanced rooted binary tree  $\tau_u$  with  $\text{deg}(u)$  leaves. We refer to the root of  $\tau_u$  as  $r_u$ . We identify each leaf of  $\tau_u$  with a child of  $u$  and denote the leaf of  $\tau_u$  that corresponds to the child  $v$  by  $c_{u,v}$ . We set the weight of edges inside  $\tau_u$  to  $\infty$ . Note that for each vertex  $u$ , the tree  $\tau_u$  has  $O(\text{deg}(u))$  vertices, and, therefore,  $T$  has  $O(n_{T'}) = O(n)$  vertices. Next, for each edge  $(u, v) \in E_{T'}$  (where we assume

that  $v$  is a child of  $u$ ), we insert the edge  $(c_{u,v}, r_v)$  in  $T$  and set  $\text{cap}_T(c_{u,v}, r_u) = \text{cap}_{T'}(u, v)$ . Finally, if  $u$  is the root of  $T'$  then we set  $r_u$  to the root of  $T$ .

It is left to show that  $T$  is a binarized Racke tree of height  $O(\log^2 n)$  and quality  $O(\log^4 n)$ . Clearly,  $T$  is a binary tree since all vertices inside each subtree  $\tau_u$  have at most two child nodes and, additionally, each vertex  $c_{u,v}$  has at most one child node (namely  $r_v$ ). Next,  $T$  has height  $O(\log^2 n)$  since  $T'$  has height  $O(\log n)$  and the subtrees  $\tau_u$  have height  $O(\log n)$ . Finally, let  $T''$  be the tree obtained from  $T$  by contracting all edges with weight  $\infty$ . We argue that  $T' = T''$ . Indeed, consider any vertex  $u \in V_{T'}$  and its subtree  $\tau_u$  in  $T$ . Then after contracting the edges in  $\tau_u$ , we are left with a subtree that only contains  $r_u$ . Furthermore, all edges between vertices of different subtrees  $\tau_u$  and  $\tau_v$ ,  $u \neq v$ , have finite weight. Therefore,  $T' = T''$ . This implies that  $T$  is binarized Racke tree for  $G$ . Since  $T'$  has quality  $O(\log^4 n)$ , the quality of  $T$  is also  $O(\log^4 n)$ . ◀

We conclude the subsection by proving Theorem 2.

**Proof of Theorem 2.** We can obtain the proof for the claim about general graphs as follows. Let  $G = (V_G, E_G, \text{cap}_G)$  be a weighted graph with  $n$  vertices. We compute a binarized Racke tree  $T = (V_T, E_T, \text{cap}_T)$  with  $O(n)$  vertices as per Lemma 29 in time  $\tilde{O}(n)$ . In  $T$ , we assign weight  $w(v) = 1$  to all vertices  $v \in V_G \cap V_T$  (i.e., to the leaves in  $T$  that correspond to vertices in  $G$ ) and weight  $w(v) = 0$  to all vertices  $v \in V_T \setminus V_G$  (i.e., to the internal nodes of  $T$  that do not correspond to any vertex in  $G$ ). Now observe that  $w(V) = n$  and thus a balanced partitioning  $V_1, \dots, V_k$  of  $T$  with  $w(V_i) \leq (1 + \epsilon) \lceil w(V)/k \rceil$  for all  $i$  corresponds to a balanced partitioning  $V'_1, \dots, V'_k$  of  $G$  with  $|V'_i| \leq (1 + \epsilon) \lceil n/k \rceil$  for all  $i$ , where  $V'_i = \{v \in V_i : w(v) = 1\}$ . Now by combining Observation 28, Proposition 26 and the fact that  $T$  has quality  $O(\log^4 n)$ , we obtain the claim.

To obtain the result about general trees  $T'$  (with unbounded degrees), we proceed similarly. We construct a binarized tree  $T$  exactly as in the proof of Lemma 29. Now, in  $T$  we set  $w(r_v) = 1$  for all root vertices of the subtrees  $\tau_v$  and we set  $w(v) = 0$  for all other vertices of the subtrees  $\tau_v$ . Similar to before, observe that  $w(V_T) = n$  and thus a balanced partitioning  $V_1, \dots, V_k$  of  $T$  with  $w(V_i) \leq (1 + \epsilon) \lceil w(V)/k \rceil$  for all  $i$  corresponds to a balanced partitioning  $V'_1, \dots, V'_k$  of  $T'$  with  $|V'_i| \leq (1 + \epsilon) \lceil n/k \rceil$  for all  $i$ , where  $V'_i = \{v \in V_{T'} : r_v \in V_i\}$ . Then by Proposition 26, this implies the proof for trees with unbounded degrees. ◀

## D.5 Extension to the Dynamic Setting

Next, we provide new dynamic algorithms in which edges are inserted and deleted from the graph. We give new algorithms for trees and for general graphs.

**Extension to Dynamic Trees.** Let us start with the case when  $T$  is a binary tree that is undergoing edge insertions and deletions. We will use Lemma 20 to make the result from Proposition 26 dynamic. However, there is a slight technical difficulty: due to edge deletions,  $T$  will become a forest and fall apart into several connected components. This becomes an issue, when an edge  $(u, v)$  is inserted for which both  $u$  and  $v$  already have parents in their respective components. In that case, we cannot immediately make  $u$  the root of  $v$  (or vice versa). Therefore, we need to find an efficient way of re-rooting the tree containing  $v$ , i.e., we need to make  $v$  the root of its component and we need to ensure that we do not have to recompute the DP solution for all vertices in the component of  $v$ . We now describe our dynamic algorithm in more detail.

First, suppose that an edge  $(u, v)$  is removed from  $T$  and assume that (before the edge deletion)  $u$  is closer to the root of  $T$  than  $v$ . Then  $T$  becomes a forest with multiple connected components. In that case, we make  $v$  the root of its component and recompute the DP

solution for  $v$  (since  $v$  does not have a parent, we only have to recompute the DP cell for  $v$ ). Furthermore, for  $u$  and all of its ancestors we recompute the DP solution as per Lemma 20.

Next, suppose an edge  $(u, v)$  is inserted, where  $u$  and  $v$  are in different connected components. Further suppose that after the edge insertion,  $u$  is the parent of  $v$ . Then we distinguish two cases whether  $v$  is the root of its component or not.

First, suppose that  $v$  is the root of its component. Then we simply insert the edge  $(u, v)$  into  $T$  and recompute the solution for  $v$  and all of its ancestors (including  $u$ ) as per Lemma 20.

Second, suppose that neither  $u$  nor  $v$  is the root of its component. Now, we first have to re-root the component containing  $v$  such that it has  $v$  as its root and such that all DP solution are valid. We do this as follows. Let  $v = v_1, \dots, v_\ell$  denote the vertices on the path from  $v$  to the root  $v_\ell$  of its component (before the edge insertion). Then we first remove all edges  $(v_\ell, v_{\ell-1}), \dots, (v_2, v_1)$  from  $T$  (in this order) as per the edge deletion routine described above. Note that after the deletions, none of the  $v_i$  has a parent and, therefore, each  $v_i$  is the root of its own component. Furthermore, by how we picked the order of the edge deletions, after the  $i$ 'th deletion we only have to recompute the DP cells for the vertices  $v_{\ell-i}$  and  $v_{\ell-i-1}$ . Now we insert the edges again but with flipped direction, i.e., we insert the edges  $(v_{\ell-1}, v_\ell), \dots, (v_1, v_2)$  (in this order). Thus,  $v = v_1$  becomes the root of the component. To insert the edges, we use the subroutine from the paragraph above, where we exploit that each  $v_i$  is the parent of its own component, which implies that the DP solutions can be updated efficiently: by how we picked the order of the edge insertions, after the  $i$ 'th edge insertion we only need to recompute the DP cells for vertices  $v_{\ell-i-1}$  and  $v_{\ell-i}$ . After the rebalancing of the component containing  $v$  is done,  $v$  has become the parent of its component and, therefore, we can use the routine from above to insert the edge  $(u, v)$ . This concludes the edge insertion procedure.

Next, when we want to output the value of the DP solution, we simply use the subroutine described in Section D.3.

We summarize the guarantees of our dynamic algorithm in the following proposition. Note that when the parameters  $\epsilon$ ,  $\bar{\epsilon}$ ,  $k$  and  $W$  are constants, the update time becomes  $\tilde{O}(h^3)$  and the query time is just  $O(1)$ . Therefore, the algorithm is very efficient for trees that have polylogarithmic or subpolynomial height in the number of vertices.

► **Proposition 30.** *Let  $\epsilon, \bar{\epsilon} > 0$  and  $k \in \mathbb{N}$ . Let  $T = (V, E, \text{cap})$  be a rooted binary tree with edge weights  $\text{cap}(e) \in W_\infty$  and vertex weights  $w(v) \in \{0, 1\}$ , that is undergoing edge insertions and deletions. Let  $h$  be an upper bound on the height of the tree  $T$  at all times. Then there exists a fully dynamic algorithm that maintains a bicriteria  $(1 + \epsilon, (1 + \bar{\epsilon})(1 + \epsilon))$ -approximation for the  $k$ -balanced partition problem on  $T$  with update time  $O(h^3 \cdot M^{2t} \log^2(W) \log(k/\epsilon) \log(M^t h \log(W)/\epsilon)/\epsilon^3)$  and query time  $M^t(k/(\epsilon\bar{\epsilon}))^{O(1/\bar{\epsilon}^2)}$ .*

**Proof.** The fact that the algorithm maintains a bicriteria  $(1 + \epsilon, (1 + \bar{\epsilon})(1 + \epsilon))$ -approximation follows immediately from Lemma 20 and the same arguments as in the proof of Proposition 26, where we argued that the approximate DP satisfies the conditions of Lemma 19.

It is left to analyze the update and query times. For the query times, note that all we do is run the subroutine from Section D.3. This subroutine runs in time  $M^t(k/(\epsilon\bar{\epsilon}))^{O(1/\bar{\epsilon}^2)}$  as we argued in the proof of Proposition 26. This proves the claim about the query time.

For the update times, let us first consider edge deletions  $(u, v)$ . In this case, we need to update the DP cell for  $v$  and the DP solutions for  $u$  and all of its ancestors. By Lemma 20, Lemma 25 and by our choice of  $p = O(h \log W/\epsilon)$ , this can be done in time  $O(h^3 \cdot M^{2t} \log^2(W) \log(k/\epsilon) \log(M^t h \log(W)/\epsilon)/\epsilon^3)$ .

Next, consider the case in which  $(u, v)$  is inserted and  $v$  is the root of its component. Then we need to recompute the DP solutions for  $v$  and all of its ancestors (including  $u$ ) which, by Lemma 20, Lemma 25 and by our choice of  $p = O(h \log W/\epsilon)$ , can be done in the time claimed in the lemma. In the case that we need to re-root the component of  $v$ , note that we have to recompute the solutions for all ancestors of  $v$ . Since the height of  $T$  is bounded by  $h$ , there are at most  $h$  such ancestors. Furthermore, we have picked the order of edge deletions such that whenever we delete or insert an edge in the re-rooting process then we only need to recompute two DP cells. Hence, in total we only need to recompute the solutions for  $O(h)$  DP cells in the re-rooting process and thus by Lemma 25, the total time for this process is  $O(h^3 \cdot M^{2t} \log^2(W) \log(k/\epsilon) \log(M^t h \log(W)/\epsilon)/\epsilon^3)$ . ◀

**Extension to Dynamic General Graphs and Non-Binary Trees.** Now suppose that our input is a dynamic (general) graph  $G$  that is undergoing edge insertions and deletions. Essentially we will solve this problem by maintaining a dynamic Racke tree and running the algorithm from Proposition 30 on top of it. However, the dynamic Racke tree from Theorem 16 is non-binary and, therefore, we start by arguing that we can maintain a binarized Racke tree dynamically in the following lemma.

► **Lemma 31.** *Let  $G = (V_G, E_G)$  be a dynamic unweighted graph with  $n$  vertices that is undergoing edge insertions and deletions. We can maintain a binarized Racke tree  $T = (V_T, E_T, \text{cap}_T)$  with  $O(n^2)$  vertices, height  $O(\log^{7/6} n)$  and quality  $n^{o(1)}$  in amortized update time  $n^{o(1)}$ . The preprocessing time is  $O(n^2)$ .*

**Proof.** Let  $T' = (V_{T'}, E_{T'}, \text{cap}_{T'})$  be the fully dynamic Racke tree for  $G$  from Theorem 16. First, note that  $T'$  has  $n_{T'} := O(n)$  vertices and height  $O(\log^{1/6} n)$ . Second, note that  $T'$  can have unbounded degree. Therefore, similar to the proof of Lemma 29, we will show how to maintain a binarized Racke tree  $T$  that has  $T'$  as its corresponding (non-binarized) Racke tree. We do so by taking  $T'$  and replacing each vertex  $u$  in  $T'$  by a balanced binary tree  $\tau_u$  with  $n_{T'}$  leaves; the internal edges of  $\tau_u$  will have weight  $\infty$  and the edges connecting subtrees  $\tau_u$  and  $\tau_v$ ,  $u \neq v$ , in  $T$  will correspond to the edges in  $T'$  and will have the same (finite) weight as in  $T'$ . We will see that by contracting all edges with weight  $\infty$  in  $T$ , we will obtain again  $T'$ . We now elaborate on this process.

During the preprocessing, we first build  $T'$ . Note that this takes time  $O(n^2)$ . Now we construct  $T$  as follows. For each vertex  $u \in V_{T'}$ , we add a balanced rooted binary tree  $\tau_u$  with  $n_{T'}$  leaves. We refer to the root of  $\tau_u$  as  $r_u$ . We identify each leaf of  $\tau_u$  with a vertex  $v \in V_{T'}$  and denote the leaf of  $\tau_u$  that corresponds to  $v$  by  $c_{u,v}$ . We set the weight of the edges inside  $\tau_u$  to  $\infty$ . Note that  $T$  has  $O(n_{T'}^2) = O(n^2)$  vertices. Next, for each edge  $(u, v) \in E_{T'}$  (where we assume that  $v$  is a child of  $u$ ), we insert the edge  $(c_{u,v}, r_v)$  in  $T$  and set  $\text{cap}_T(c_{u,v}, r_u) = \text{cap}_{T'}(u, v)$ . Finally, if  $u$  is the root of  $T'$  then we set  $r_u$  to the root of  $T$ .

Next, suppose that  $G$  is changed due to an edge insertion or deletion. Then we first update the tree  $T'$  via the algorithm from Theorem 16. Now, whenever an edge  $(u, v)$  is inserted (deleted) in  $T'$ , we insert (delete) the edge  $(c_{u,v}, r_v)$  into (from)  $T$ . Each of these insertions and deletions in  $T$  can be done in time  $O(1)$ . Since it takes amortized time  $n^{o(1)}$  to update  $T'$  (via Theorem 16), the total update time is  $n^{o(1)}$ .

It is left to show that  $T$  is a binarized Racke tree of height  $O(\log^{7/6} n)$  and quality  $n^{o(1)}$ . Clearly,  $T$  is a binary tree since all vertices inside each subtree  $\tau_u$  have at most two child nodes and, additionally, each vertex  $c_{u,v}$  has at most one child node (namely  $r_v$ ). Next,  $T$  has height  $O(\log^{7/6} n)$  since  $T'$  has height  $O(\log^{1/6} n)$  and the subtrees  $\tau_u$  have height  $O(\log n)$ . Finally, let  $T''$  be the tree obtained from  $T$  by contracting all edges with weight  $\infty$ . We argue that  $T' = T''$ . Indeed, consider any vertex  $u \in V_{T'}$  and its subtree  $\tau_u$  in  $T$ .

Then after contracting the edges in  $\tau_u$ , we are left with a subtree that only contains  $r_u$ . Furthermore, all edges between vertices of different subtrees  $\tau_u$  and  $\tau_v$ ,  $u \neq v$ , have finite weight. Therefore,  $T' = T''$ . This implies that  $T$  is binarized Racke tree for  $G$ . Since  $T'$  has quality  $n^{o(1)}$ , the quality of  $T$  is also  $n^{o(1)}$ . ◀

Given the lemma above, our dynamic algorithm for dynamic general graphs  $G$  works as follows. We maintain the dynamic binarized Racke tree  $T$  as per Lemma 31 on our input graph  $G$ , i.e., whenever an edge is inserted or deleted in  $G$ , we update the data structure from the lemma as well. Note that this causes edge insertions and deletions in  $T$  as well. As before, we set the vertex weights in  $T$  such that  $w(v) = 1$  if  $v$  corresponds to a vertex in  $G$  and  $w(v) = 0$  if  $v$  is an internal node of  $T$  that does not correspond to any vertex in  $G$ . Furthermore, we run our dynamic algorithm from Proposition 30 for binary trees on  $T$ . In particular, whenever  $T$  gets updated, we also update the DP solution as per Proposition 30. We also use the same query procedure as in the proposition.

We conclude the subsection by proving Theorem 3.

**Proof of Theorem 3.** We prove the result for general graphs first. Since the dynamic binarized Racke tree  $T$  that we maintain has quality  $n^{o(1)}$ , the same argumentation as in the proof of Theorem 2 implies that we maintain a bicriteria  $(n^{o(1)}, (1 + \bar{\epsilon})(1 + \epsilon))$ -approximation for  $G$ . Since by Lemma 31 we can maintain  $T$  with amortized update time  $n^{o(1)}$ , the amortized number of edge insertions and deletions into  $T$  is  $n^{o(1)}$  per update operation. Since  $T$  has height  $O(\log^{7/6} n)$  and by Proposition 30, the total amortized update time  $n^{o(1)}$ . This implies the claim about dynamic general graphs.

To obtain our result for non-binary trees, we can proceed similar to above. Consider a non-binary  $T'$  that is undergoing edge insertions and deletions. We can maintain the same data structure as in the proof of Lemma 31 to obtain a binary tree  $T$  with  $O(n^2)$  vertices with worst-case update time  $O(1)$ . Now we assign weight  $w(r_v) = 1$  to all vertices  $r_v$  that are roots of the subtrees  $\tau_v$  in  $T$  and weight  $w(v) = 0$  to all other vertices of the subtrees  $\tau_v$ . By the same arguments as in the proof of Theorem 2, we obtain the claim. ◀

## E Simultaneous Source Location

In this section, we provide efficient algorithms for the simultaneous source location problem as studied by Andreev et al. [4]. Recall that in this problem, the input consists of a graph  $G = (V, E, \text{cap}, d)$  with a capacity function  $\text{cap}: E \rightarrow W_\infty$  on the edges and a demand function  $d: V \rightarrow W_\infty$  on the vertices of the graph. The goal is to select a minimum set  $S \subseteq V$  of *sources* that can simultaneously supply all vertex demands. More concretely, a set of sources  $S$  is *feasible* if there exists a flow from the vertices in  $S$  that supplies demand  $d(v)$  to all vertices  $v \in V$  and that does not violate the capacity constraints on the edges. Here, we assume that each source vertex can potentially send an infinite amount of flow that is only constrained by the edge capacities. The objective is to find a feasible set of sources of minimum size.

Next, we summarize our main results for the simultaneous source location problem. First, we introduce our notion of bicriteria approximation. Let  $S^*$  be the optimal solution for the simultaneous source location problem. Then we say that a solution  $S$  is a bicriteria  $(\alpha, \beta)$ -approximate solution if  $|S| \leq \alpha |S^*|$  and if  $S$  is a feasible set of sources after all edge capacities are increased by a factor  $\beta$ .

The following theorem summarizes our main result for static algorithms.



► **Theorem 4.** *Let  $\epsilon > 0$ . Let  $G = (V, E, \text{cap}, d)$  be an undirected weighted graph with  $n$  vertices and  $m$  edges. Then for the simultaneous source location problem we can compute:*

- *A  $(1 + \epsilon, O(\log^4(n)))$ -approximation in time<sup>12</sup>  $\tilde{O}(\frac{1}{\epsilon^2}m)$ .*
- *A  $(1 + \epsilon, 1)$ -approximation in time  $\tilde{O}(\frac{1}{\epsilon^2}h^2 \cdot n)$  if  $G$  is a tree of height  $h$ .*

Next, we turn to our dynamic algorithms which support the following update operations:

- *SetDemand( $v, d$ ): updates the demand of vertex  $v$  to  $d(v) = d$ ,*
- *SetCapacity( $(u, v), c$ ): updates the capacity of the edge  $(u, v)$  to  $\text{cap}(u, v) = c$ ,*
- *Remove( $u, v$ ): removes the edge  $(u, v)$  from the graph,*
- *Insert( $(u, v), c$ ): inserts the edge  $(u, v)$  into the graph with capacity  $\text{cap}(u, v) = c$ .*

The next theorem summarizes our main results for dynamic algorithms.

► **Theorem 5.** *Let  $\epsilon > 0$ . Let  $G = (V, E, \text{cap}, d)$  be a graph with  $n$  vertices and  $m$  edges that is undergoing the update operations given above. Then for the simultaneous source location problem we can maintain:*

- *A  $(1 + \epsilon, n^{o(1)})$ -approximation with amortized update time  $n^{o(1)}/\epsilon^2$  and preprocessing time  $O(n^2/\epsilon^2)$  if all edge capacities are 1.*
- *A  $(1 + \epsilon, O(\log^4(n)))$ -approximation with worst-case update time  $\tilde{O}(1/\epsilon^2)$  and preprocessing time  $\tilde{O}(m)$  if we only allow the update operation SetDemand( $v, d$ ).*
- *A  $(1 + \epsilon, O(\log^2(n) \log \log(n)))$ -approximation with worst-case update time  $\tilde{O}(1/\epsilon^2)$  and preprocessing time  $\text{poly}(n)$  if we only allow the update operation SetDemand( $v, d$ ).*
- *A  $(1 + \epsilon, 1)$ -approximate solution with worst-case update time  $\tilde{O}(h^3/\epsilon^2)$  and preprocessing time  $O(n^2/\epsilon^2)$  if  $G$  is a tree of height  $h$ .*

We note that in our static and dynamic algorithms, we can output the corresponding solutions similarly to what we described after Proposition 12 for knapsack.

We start by presenting an exact DP for the special case of binary trees in Section E.1 and then present an approximate DP in Section E.2. After that, we generalize the result from binary trees to general graphs in Section E.3 and then also to the fully dynamic setting in Section E.4.

## E.1 The Exact DP

We consider the special case of the simultaneous source location problem on *binary* trees and provide a DP that solves this problem exactly. We let  $T = (V, E, \text{cap}, d)$  denote the rooted binary tree with root  $r$  that we obtain as input. Additionally, we assume that for each vertex  $v \in V$  we obtain as input whether we are allowed to make  $v$  a source or not; note that this only generalizes the problem (as in the original problem all vertices can be made sources). Later in Section E.3, this generalization will be helpful when we apply Räcke trees because then we only want to allow leaves to act as sources.

### E.1.1 DP Definition

We now define our exact DP. We will also discuss its relationship with the DP by Andreev et al. [4] and why we did not use the DP of Andreev et al. Given a vertex  $v$  and a value  $x \in \mathbb{R}$ , we denote by  $\text{DP}(v, x)$  the minimum number of sources to place in  $T_v$  such that when  $v$  receives flow at most  $x$  from its parent then all demands in  $T_v$  can be satisfied. We note that  $x$  can take positive and negative values: for  $x \geq 0$  this corresponds to the setting in

<sup>12</sup>We write  $\tilde{O}(f(n, \epsilon, W))$  to denote running times of the form  $f(n, \epsilon, W) \cdot \text{polylog}(n, \epsilon, \log W)$ .

which flow is sent from the parent of  $v$  into  $T_v$  and for  $x < 0$  this corresponds to the setting in which flow is sent from  $T_v$  towards the parent of  $v$ . We further follow the convention that when the demands in  $T_v$  cannot be satisfied when  $v$  receives flow  $x$  from its parent, then we set  $\text{DP}(v, x) = \infty$ .

Observe that this DP has rows  $\mathcal{I} = V$  and columns  $\mathcal{J} = \mathbb{R}$ . We will store the rows  $\text{DP}(v, \cdot)$  using our data structure from Section 2 using monotone piecewise constant functions. Next, we observe that each  $\text{DP}(v, \cdot)$  is monotonically decreasing. Hence, the DP satisfies Property (1) of Definition 8.

► **Observation 32.** *The function  $\text{DP}(v, \cdot): \mathbb{R} \rightarrow [n + 1] \cup \{\infty\}$  is monotonically decreasing.*

**Proof.** This follows immediately from the definition of  $\text{DP}(v, x)$ : Consider  $x, x' \in \mathbb{R}$  with  $x \leq x'$ . Then any solution in which  $T_v$  receives flow at most  $x$  from the parent of  $v$  is also feasible when  $T_v$  receives flow at most  $x'$  from the parent of  $v$ . Therefore,  $\text{DP}(v, x) \geq \text{DP}(v, x')$ , which finishes the proof. ◀

Observe that the global solution for the simultaneous source location problem on  $T$  can be obtained by evaluating  $\text{DP}(r, 0)$ , where  $r$  is the root of  $T$ : First,  $r$  has no parent and, therefore, it must be a source itself or have its demand satisfied by its children; this explains the choice of  $x = 0$ . Furthermore, (by definition)  $\text{DP}(r, 0)$  is the minimum number of sources that we need to satisfy all demands in  $T_r = T$  and, thus, the flow that we obtain is feasible. We conclude that  $\text{DP}(r, 0)$  gives the global optimum solution.

**Relationship to the approach by Andreev et al. [4].** Next, let us elaborate on the relationship of our DP and the function  $f$  used by Andreev et al. [4]. In [4], the function  $f$  computed by a dynamic program is defined as follows. Given a vertex  $v$  and an integer  $i \in \mathbb{N}$ , Andreev et al. define a function  $f(v, i)$  that denotes the *minimum amount of flow* that  $v$  needs to receive from its parent if all demands in  $T_v$  need to be satisfied and if we can place  $i$  sources in the subtree  $T_v$ . Similar to above,  $f(v, i)$  can take positive and negative values: if the demand in  $T_v$  can only be satisfied by receiving flow from the parent, then  $f(v, i)$  is positive; if the demand in  $T_v$  is already satisfied by the sources in the subtree  $T_v$ , then it is possible that  $v$  can send flow to its parent and  $f(v, i)$  is negative. It is not hard to see that the function  $f(v, i)$  is monotonically decreasing in  $i$ .<sup>13</sup>

Now consider  $f(v, \cdot): \mathbb{N} \rightarrow \mathbb{R}$  as a function and consider its “inverse”<sup>14</sup> function  $f^{-1}(v, \cdot): \mathbb{R} \rightarrow \mathbb{N}$ , where  $f^{-1}$  is defined on the whole set of real numbers (including negative numbers). That is,  $f^{-1}(v, x)$  denotes the minimum number of sources that we need to place in  $T_v$  such that the demand that  $v$  requires from its parent is at most  $x$ . But this was exactly the definition of  $\text{DP}(v, x)$ . Thus,  $\text{DP}(v, x) = f^{-1}(v, x)$  for all  $v$  and  $x$ .

*Why Did We Not Use  $f$ ?* In [4] it is shown how the function  $f$  can be computed in polynomial time by a bottom-up dynamic program using just a few case distinctions and a  $(\min, +)$ -convolution in each DP cell. Thus, one might wonder why we picked  $\text{DP}(v, \cdot) = f^{-1}(v, \cdot)$  and not  $f$  for our DP? Indeed, it seems quite natural to interpret the function  $f$  as a monotone piecewise constant function and to use it for our dynamic program.

<sup>13</sup>This follows immediately from the definition of  $f$  and the fact that by adding more sources to a subtree  $T_v$ , the amount of flow that  $T_v$  needs to receive from the parent of  $v$  only decreases.

<sup>14</sup>We note that, formally,  $f(v, \cdot)$  has no inverse since it is possible that multiple values map to the same number, i.e.,  $f(v, i) = f(v, i')$  for  $i \neq i'$ . Thus, formally, we set  $f^{-1}(v, x) = \min\{i: f(v, i) \leq x\}$ , where we follow the convention  $\min\{\emptyset\} = \infty$ . Then we interpret  $f^{-1}(v, \cdot)$  as a piecewise constant function from  $\mathbb{R}$  to  $[n + 1]$ .

While for the case of exact computations this is possible, we now sketch why this appears unhandy for the approximate case later.

Suppose that we used the function  $f$  in our approximate computations. To obtain efficient approximation algorithms, we will have to ensure that  $f$  has only few pieces and our main way to achieve this is by rounding  $f$  as per Lemma 6. However, this becomes tricky because the function values of  $f$  are positive *and negative*. In the following, it will be illustrative to think of positive function values for  $f$  as vertex demands that need to be satisfied and of negative values for  $f$  as available edge capacities. The main issue is that since the function values of  $f$  are positive and negative, it is not clear how we should perform the rounding: if we rounded positive and negative values up (towards  $+\infty$ ) then this would correspond to *increasing* the vertex demands while at the same time *decreasing* the edge capacities; however, this could render some feasible solutions (in the exact computation) infeasible (in the rounded computation). On the other hand, it is conceivable that by always rounding  $f$  *down* (towards  $-\infty$ ), we would essentially decrease the vertex demands while increasing the edge capacities. Potentially, this approach could work when we are allowed to violate the edge capacities by a  $(1 + \epsilon)$ -factor. However, even if we did that, we would have another issue: to only use a small number of pieces for representing  $f$ , we would have to use different rounding mechanisms for those function values in  $[-1, 1]$  and those in  $[-W, W] \setminus [-1, 1]$ , where  $W$  is the largest edge capacity. Indeed, if we rounded the values of  $f$  to powers of  $(1 + \delta)^j$  then there are only  $O(\log_{1+\delta}(W))$  function values in  $[-W, W] \setminus [-1, 1]$  but there are infinitely many function values in  $[-1, 1]$ . Similarly, if we rounded to multiples of  $\delta$  then there are only  $O(1/\delta)$  function values in  $[-1, 1]$  but this would lead to  $O(W/\delta)$  function values in  $[-W, W] \setminus [-1, 1]$ . In both cases, our functions would have too many pieces and, thus, one would have to pick a rounding function which provides a tradeoff between these two cases. Furthermore, we would have to find an analysis that shows that this “more involved” rounding function does not introduce too much error.

Note that in the above discussion, all of the issues come from the fact that  $f(v, \cdot)$  can also take negative values. On the other hand, our DP (which is  $f^{-1}(v, \cdot)$ ) only takes non-negative function values and, therefore, we avoid all of the above complications because we can use the standard rounding function  $\lceil \cdot \rceil_{1+\delta}$  that rounds to powers of  $1 + \delta$ . Thus, we bypass all of the issues above.

Our approach also has the positive side effects that instead of getting factors of  $\text{polylog}(W)$  in our running times, we only get factors of  $\text{polylog}(n)$  because the codomain of our monotone piecewise constant functions became  $[n + 1]$  rather than some potentially large interval  $[-W, W]$ .

## E.1.2 Computing the DP

Now we describe the exact computation of our DP. This will reveal the procedures  $\mathcal{P}_i$  from Definition 8.

Let  $v \in V$  be any vertex in  $T$ . We describe how to compute  $\text{DP}(v, \cdot)$  efficiently assuming that we have already computed the solutions for the children of  $v$  (if they exist). Recall that for each vertex  $v$  we also obtain as input, whether  $v$  can be used as a source or not. In our following case distinctions, whenever we consider the case that  $v$  is used as a source, we will implicitly condition on the fact that it is also possible to use  $v$  as source; if  $v$  cannot be used as a source, we simply skip this case.

In the construction for each DP cell  $\text{DP}(v, \cdot)$  for a vertex  $v$  with parent  $p$ , we will additionally ensure that we do not violate the capacity of the edge  $(p, v)$  when  $x$  is very small or very large. More concretely, we will ensure that  $\text{DP}(v, \cdot)$  satisfies the additional property

that  $DP(v, x) = \infty$  for  $x < -\text{cap}(p, v)$  and  $DP(v, x) = DP(v, \text{cap}(p, v))$  for all  $x > \text{cap}(p, v)$ . We will denote this property as the *feasible capacity property*.

**Case 1:  $v$  is a leaf.** Suppose that  $v$  is a leaf. We initialize  $DP(v, \cdot)$  as the function which takes value  $\infty$  on all of  $\mathbb{R}$ . In the following, we add at most two pieces to  $DP(v, \cdot)$  depending on whether  $v$  can be used as a source or not.

First, suppose  $v$  can be used as a source. Then we can send flow up to  $\text{cap}(p, v)$  to the parent  $p$  of  $v$ . Furthermore, since  $v$  is a leaf, there is exactly one source in  $T_v$ . Thus, we update  $DP(v, \cdot)$  and set  $DP(v, x) = 1$  for all  $x \geq -\text{cap}(p, v)$ . This adds one piece to  $DP(v, \cdot)$ .

Second, suppose  $v$  is not a source. Then if  $x \geq d(v)$  and  $\text{cap}(p, v) \geq d(v)$ ,  $v$  can receive all of its demand  $d(v)$  from its parent and the flow is feasible because we do not exceed the capacity of the edge  $(p, v)$ . Therefore, if  $\text{cap}(p, v) \geq d(v)$ , then we update  $DP(v, \cdot)$  again and add the piece with  $DP(v, x) = 0$  for all  $x \geq d(v)$ . If  $x \geq d(v)$  but  $d(v) > \text{cap}(p, v)$  then we do nothing because the parent of  $v$  cannot satisfy the demand of  $v$ .

Observe that  $DP(v, \cdot)$  is a monotonically decreasing function with at most three pieces. Furthermore, it clearly satisfies the feasible capacity property and Property (3) of Definition 8.

**Case 2:  $v$  is not a leaf.** Suppose that  $v$  is not a leaf and that  $v$  has children  $v_1$  and  $v_2$ , as well as a parent  $p$ . Recall that we assume that we have already computed the DP entries  $DP(v_1, \cdot)$  and  $DP(v_2, \cdot)$  for both children of  $v$ . We now show how to compute two DP solutions  $DP_A(v, \cdot)$  and  $DP_B(v, \cdot)$  depending on whether  $v$  is a source (in Case A) or not (in Case B). Then, if  $v$  can be used as a source, we set

$$DP(v, \cdot) = \min\{DP_A(v, \cdot), DP_B(v, \cdot)\},$$

where we compute the min-operation via Lemma 6. If  $v$  cannot be used as a source, we set  $DP(v, \cdot) = DP_B(v, \cdot)$ .

*Case A:* Suppose  $v$  is used as a source. We initialize  $DP_A(v, \cdot)$  as the function which takes value  $\infty$  on all of  $\mathbb{R}$ . Now, since  $v$  can be used as a source,  $v$  can send flow  $\text{cap}(p, v)$  to its parent and flow  $\text{cap}(v, v_1)$  and  $\text{cap}(v, v_2)$  to its children. Therefore, for  $x \geq -\text{cap}(p, v)$ , the number of sources in  $DP_A(v, x)$  is 1 (since  $v$  is a source) *plus* the number of sources that we require in  $T_{v_1}$  when  $v_1$  can receive flow  $\text{cap}(v, v_1)$  from its parent  $v$  *plus* the same quantity for  $T_{v_2}$ . Thus, it suffices to set

$$DP_A(v, x) = 1 + DP(v_1, \text{cap}(v, v_1)) + DP(v_2, \text{cap}(v, v_2))$$

for all  $x \geq -\text{cap}(p, v)$ . Note that here we exploited that the functions  $DP(v_1, \cdot)$  and  $DP(v_2, \cdot)$  are monotonically decreasing and that both of them satisfy the feasible capacity property. We conclude that in this case  $DP_A(v, \cdot)$  is a monotonically decreasing function with two pieces.

*Case B:* Suppose that  $v$  is not used as a source. We initialize  $DP_B(v, \cdot)$  as the function which takes value  $\infty$  on all of  $\mathbb{R}$ . To compute the value of  $DP_B(v, x)$ , we need to obtain the minimum number of sources such that  $v$  receives flow at most  $x$  from its parent and such that all demands in  $T_v$  are satisfied. Since  $v$  is not a source, its demand  $d(v)$  must be satisfied either by its parent  $p$  or by its children (or a combination of them). Therefore, to obtain that we have to pick the children solutions  $DP(v_1, x_1)$  and  $DP(v_2, x_2)$  such that  $d(v) \leq x - x_1 - x_2$ .

Since we did not make  $v$  a source, the number of sources in  $DP_B(v, x)$  is the number of sources that we need to place in the subtrees  $T_{v_1}$  and  $T_{v_2}$ . Thus, we get

$$DP_B(v, x) = \min_{x_1 \in \mathbb{R}} \{DP(v_1, x_1) + DP(v_2, x - x_1 - d(v))\},$$

where we used that  $x_2 \leq x - x_1 - d(v)$  and by monotonicity of  $\text{DP}(v_2, \cdot)$  we minimize the number of sources in  $T_{v_2}$  if we consider  $x_2 = x - x_1 - d(v)$ . Here, the flows that we computed for  $\text{DP}_B(v, x)$  are set feasible because the solutions  $\text{DP}(v_i, \cdot)$  satisfy the feasible capacity property and therefore we do not violate the edge constraints to the children.

Since the above equality holds for all values of  $x$ ,  $\text{DP}_B(v, \cdot)$  corresponds to a shifted  $(\min, +)$ -convolution of two monotonically decreasing functions. More concretely, via Lemma 6 we can first compute the shifted function  $\text{DP}(v_2, \cdot - d(v))$  and then we can set

$$\text{DP}_B(v, \cdot) = \text{DP}(v_1, \cdot) \oplus \text{DP}(v_2, \cdot - d(v)),$$

which we compute via Lemma 7.

Finally, as a postprocessing step, we set  $\text{DP}(v, \cdot) = \min\{\text{DP}_A(v, \cdot), \text{DP}_B(v, \cdot)\}$  if  $v$  can be used as a source and  $\text{DP}(v, \cdot) = \text{DP}_B(v, \cdot)$  otherwise, as we already mentioned above. But we also need to ensure that  $\text{DP}(v, \cdot)$  satisfies the feasible capacity property. Therefore, we set  $\text{DP}(v, x) = \infty$  for  $x < -\text{cap}(p, v)$  and we set  $\text{DP}(v, \cdot) = \text{DP}(v, \text{cap}(p, v))$  for  $x > \text{cap}(p, v)$ . Observe that these changes to  $\text{DP}(v, \cdot)$  can be done in time linear in the number of pieces of  $\text{DP}(v, \cdot)$ .

**Properties of the DP.** Observe that in the DP above for each vertex  $v$  we only required the DP solutions for its children  $v_1$  and  $v_2$ . Hence, our dependency graph is given by our input tree  $T$  where all edges are directed towards the root. This implies that every node in the dependency graph can only reach those nodes on a path to the root and thus Property (2) of Definition 8 is satisfied with  $h$  being the height of  $T$ . Additionally, one can verify that above all operations also satisfy Property (3) of Definition 8. Finally, observe that in each step we only used a constant number of operations from Lemma 6 and at most one  $(\min, +)$ -convolution from Lemma 7.

## E.2 The Approximate DP

Now we explain how we solve the above DP more efficiently by computing approximate solutions  $\text{ADP}(v, \cdot)$ . This will reveal the procedures  $\tilde{\mathcal{P}}_i$  from Definition 8.

In our approximation algorithm, we do everything exactly as above except that we replace each exact solution  $\text{DP}(v, \cdot)$  with the approximate solution  $\text{ADP}(v, \cdot)$ . Then we add a postprocessing step in which we round  $\text{ADP}(v, \cdot)$ , i.e., we set

$$\text{ADP}(v, \cdot) = \lceil \text{ADP}(v, \cdot) \rceil_{1+\delta} \tag{12}$$

for a parameter  $\delta > 0$  that we will set later.

Note that all of our operations are exact except the rounding step which loses a factor of  $\alpha = 1 + \delta$ . Thus, Property (4a) of Definition 8 is satisfied. Additionally, observe that in each step we only used a constant number of operations from Lemma 6 and at most one  $(\min, +)$ -convolution from Lemma 7. This implies that Property (4b) is satisfied. Furthermore, all functions we consider are monotone and our rounding step ensures that each row  $\text{ADP}(v, \cdot)$  has at most  $p = O(\log_{1+\delta} n)$  pieces. Hence, Property (4c) is also satisfied.

This implies that the DP is  $(h, 1 + \delta, O(\log_{1+\delta}(n)))$ -well-behaved. By applying Theorem 9 with  $\delta = \ln(1 + \epsilon)/(h + 1)$ , we obtain the following proposition which shows that on binary trees, the approximation algorithm computes a bicriteria  $(1 + \epsilon, 1)$ -approximate solution. We note that for constant  $\epsilon$ , the running time essentially becomes  $\tilde{O}(n \cdot h^2)$ , where  $h$  is the height of the tree. Thus, for trees of height  $\tilde{O}(1)$ , we obtain a near-linear running time.

► **Proposition 33.** *Let  $\epsilon > 0$ . The approximation algorithm computes a bicriteria  $(1 + \epsilon, 1)$ -approximate solution for the simultaneous source location problem on binary trees in time  $O(n \cdot (h \log(n)/\epsilon)^2 \log(h \log(n)/\epsilon))$ , where  $h$  is the height of the tree.*

### E.3 Extension to General Graphs (Proof of Theorem 4)

We prove Theorem 4 by giving reductions to the binary setting.

First, suppose that  $G$  is a tree with potentially unbounded degree. Then we turn  $G$  into a binary tree  $T$  using the same construction as in the proof of Lemma 29. That is, we replace each vertex  $u$  in  $G$  by a balanced binary tree  $\tau_u$  with  $\deg(u)$  leaves  $c_{u,v_1}, \dots, c_{u,v_{\deg(u)}}$ , where the  $v_i$  are the children of  $u$  in  $G$ ; the internal edges of  $\tau_u$  have capacity  $\infty$  and we denote the root of each  $\tau_u$  by  $r_u$ . Furthermore, for each edge  $(u, v)$  in  $G$ , we insert the edge  $(c_{u,v}, r_v)$  into  $T$  with capacity  $\text{cap}(c_{u,v}, r_v) = \text{cap}(u, v)$ . By the same arguments as in the proof of Lemma 29,  $T$  has  $O(n)$  vertices and height  $O(h \log n)$ , where  $h$  is the height of  $G$ . It is straight-forward to see that with this construction, there exists a flow from  $u$  to  $v$  in  $G$  if and only if there exists a flow from  $r_u$  to  $r_v$  in  $T$ . Now, in  $T$  we have already set the edge capacities and it remains to set the vertex demands. For each vertex  $r_u$  in  $T$ , we set  $d(r_u) = d(u)$ , and for all other vertices  $v$  in  $T$ , we set  $d(v) = 0$ . Furthermore, in our instance of the simultaneous source location problem we set that each vertex  $r_u$  can be picked as a source and none of the other vertices in  $T$  can be picked as a source. Note that there exists a one-to-one correspondence between sources in  $G$  and sources in  $T$ . Together with our observation for flows above, this means that solving the simultaneous source location problem on  $T$  gives a solution for  $G$ .

To obtain the bicriteria  $(1 + \epsilon, 1)$ -approximation result for trees, we apply the approximation algorithm from Proposition 33 on  $T$ .

Finally, to obtain the  $(1 + \epsilon, O(\log^4 n))$ -approximate solution for a general graph  $G$ , we proceed as follows. We build the binarized Racke tree  $T$  for  $G$  as per Lemma 29 and recall that  $T$  has quality  $q = O(\log^4 n)$  and height  $\tilde{O}(1)$ . In  $T$ , we set the bits to indicate that all leaves can be used as sources but none of the other vertices might be used as a source. We apply the approximation algorithm from Proposition 33 on  $T$  to obtain a  $(1 + \epsilon, 1)$ -approximate solution on  $T$  in time  $\tilde{O}(n)$ . Now let us point out that the Racke tree from Theorem 15 (and, therefore, also the binarized Racke tree from Lemma 29) is also a *tree flow sparsifier*. That is, if there exists a feasible flow  $F$  in  $G$ , then there exists a flow of the same value between the corresponding leaves in  $T$ . Additionally, for any feasible flow  $F$  with value  $v$  between leaves in  $T$ , there exists a feasible flow with value  $\frac{1}{q}v$  between the corresponding vertices in  $G$ . Therefore, if we are allowed to exceed the edge capacities in  $G$  by a factor of  $q = O(\log^4 n)$ , the flow that we compute in  $T$  is feasible in  $G$ . This gives that we can compute a  $(1 + \epsilon, O(\log^4 n))$ -approximate solution in time  $\tilde{O}(n)$ .

### E.4 Extension to the Dynamic Setting (Proof of Theorem 5)

To prove Theorem 5, we first consider the special case of dynamic *binary* trees (which is not mentioned in the theorem). We show that for binary trees we can maintain a bicriteria  $(1 + \epsilon, 1)$ -approximate solution with worst-case update time  $\tilde{O}(h^3/\epsilon^2)$ , where  $h$  is an upper bound on the height of the tree. Then we show that the results of the theorem can be derived from this result.

Consider a dynamic *binary* tree on which we maintain the approximate DP from Section E.2. We will exploit that  $T$  and the dependency tree of our DP coincide. Hence, an update in  $T$  will trigger the same update in the dependency tree. Observe that the update operation  $\text{SetDemand}(v, d)$  triggers a change to  $\text{ADP}(v, \cdot)$ . Then we can recompute the global approximate DP table using Theorem 10. Since the DP is well-behaved, the tree has height at most  $h$  and since we set  $\delta = O(h/\epsilon)$ , the theorem implies that we need time  $\tilde{O}(h^3/\epsilon^2)$  to recompute the ADP solution. Similarly, for  $\text{SetCapacity}((u, v), c)$  we can again update the

rows  $\text{ADP}(u, \cdot)$  and  $\text{ADP}(v, \cdot)$  and we update the entire DP table using Theorem 10. By the same arguments as above, this takes time  $\tilde{O}(h^3/\epsilon^2)$ . For  $\text{Remove}(u, v)$ , we remove the edge  $(u, v)$  from  $T$  and by the same reasoning as before we get update time  $\tilde{O}(h^3/\epsilon^2)$ . Finally, consider  $\text{Insert}((u, v), c)$ , where we assume that  $v$  becomes the child of  $u$ . Then we might have the issue that before the update,  $v$  is not the root of its connected component. To mitigate this issue, we run the same re-rooting procedure as described in Section D.5. As described in Section D.5, this will only recompute the solutions of  $O(h)$  DP cells and thus we again have a total update time of  $\tilde{O}(h^3/\epsilon^2)$ .

Next, we prove the results from Theorem 5. First, consider the case in which all edge capacities are set to 1 and where we want to obtain a bicriteria  $(1 + \epsilon, n^{o(1)})$ -approximate solution with amortized update time  $n^{o(1)}/\epsilon^2$  and preprocessing time  $O(n^2)$ . Let  $G$  be the dynamic input graph. We maintain the dynamic binarized Räcke tree  $T$  for  $G$  as per Lemma 31 and remark that the dynamic Räcke tree from Theorem 16 is also a tree flow sparsifier. We note that any update to  $G$  triggers an update operation on  $T$  that requires amortized update time  $n^{o(1)}$ . On  $T$ , we allow the leaves to act as sources but no other vertices. Furthermore, we set the demands of the leaves in  $T$  to the demands of the corresponding vertices in  $G$ ; all other vertices have demand 0. Now we use the data structure for binary trees from the previous paragraph to maintain a dynamic bicriteria  $(1 + \epsilon, 1)$ -approximate solution on  $T$ . That is, when a vertex demand changes in  $G$ , we update the corresponding vertex demand in  $T$ . When an edge is inserted or deleted in  $T$  due to the subroutine from Lemma 31, then we update the data structure from the previous paragraph that maintains the DP solution on  $T$ . By the same argumentation as in Section E.3, we obtain that since  $T$  has quality  $n^{o(1)}$ , if we can exceed the edge capacities in  $G$  by a  $n^{o(1)}$  factor then any feasible flow in  $T$  is also feasible in  $G$ . This implies the result claimed in the theorem.

If  $G$  is a tree of height  $h$  but (potentially) with unbounded degrees, we can maintain a bicriteria  $(1 + \epsilon, 1)$ -approximate solution with worst-case update time  $\tilde{O}(h^3/\epsilon^2)$  and preprocessing time  $O(n^2/\epsilon^2)$  similar to above. That is, we transform  $G$  into a binary tree using the same procedure that we use in the proof of Lemma 31, where we replace each vertex  $u$  by a subtree  $\tau_u$  with root  $r_u$ . Similar to what we argued in Section E.3, we only allow the vertices  $r_u$  as roots in  $T$  and obtain any flow in  $T$  corresponds to a flow in  $G$ . Then by applying the dynamic data structure for binary trees on  $T$ , we obtain the result.

Finally, let us consider the case in which we wish to obtain bicriteria approximation algorithms when we only allow the update operations  $\text{SetDemand}(v, d)$ . In this case, we observe that the underlying graph is static, since only the vertex demands change. Therefore, for our input graph  $G$ , we can build the Räcke tree from Theorem 15 which is also a tree flow sparsifier and we consider its binarized version  $T$  as per Lemma 31. Note that building this tree with quality  $\tilde{O}(\log^4 n)$  takes time  $\tilde{O}(m)$ . Given such a static Räcke tree, we can use our dynamic data structure for binary trees from above to support the operations  $\text{SetDemand}(v, d)$  on  $T$ . Since  $T$  has height  $\tilde{O}(1)$ , we obtain the result with the bicriteria  $(1 + \epsilon, O(\log^4 n))$ -approximation. To obtain the bicriteria  $(1 + \epsilon, O(\log^2 n \log \log n))$ -approximation we do exactly the same as above, but instead of using the Räcke tree from Theorem 15, we use the Räcke tree from Harrelson, Hildrum and Rao [37] which can be used as a tree flow sparsifier. As it has quality  $O(\log^2 n \log \log n)$  and can be constructed in time  $\text{poly}(n)$ , we obtain the result.

## F Recourse Bounds

In this section we discuss the recourse bounds we derive. To motivate these lower bounds, let us note that “classic” dynamic algorithms with polylogarithmic update time maintain a single explicit solution in memory; this is desirable in many practical scenarios. However, some dynamic algorithms (like our DP algorithms above) only return the *value* of an approximate solution in polylogarithmic time, which is sometimes referred to as *implicit*. To understand whether for our problems implicit solutions are necessary, we consider algorithms which maintain multiple explicit solutions, of which only one has to be feasible. We believe that this is an interesting setting to look at, as it essentially interpolates between the two scenarios above. If even algorithms with multiple solutions must have high recourse, this suggests that implicit solutions are somehow inevitable. We show below that for fully dynamic knapsack and fully dynamic  $k$ -balanced partitioning the latter is the case.

Here, we consider dynamic algorithms over inputs that are undergoing insertions and deletions via an update operation. The algorithms are allowed to maintain multiple explicit solutions and must ensure that after every time step, there exists a solution with certain guarantees while minimizing the recourse for updating the solutions. More concretely, we consider algorithms which explicitly maintain  $s$  solutions  $S_1^{(t)}, \dots, S_s^{(t)}$  for each time step  $t$ . Here, we assume that after each time step a single update operation is performed, after which an algorithm can make changes to its solutions. We say that an algorithm maintains an  $\alpha$ -*approximate solution* if for each time step  $t$ , there exists an index  $i = i(t)$  such that  $S_i^{(t)}$  is a feasible and  $\alpha$ -approximate solution for the problem we study.

Observe that in this setting, the algorithm might have much lower recourse, since for each time  $t$  it may pick a different solution. Thus, it may not have to update any of the solutions significantly after the update operations. Further note that this notion of ensuring that at each time step there exists a feasible solution is somewhat reminiscent of list decoding in coding theory, where the decoder can output a list of messages and only has to ensure that the correct message is contained in that list.

Measuring the recourse will be problem-specific, based on how the solutions for the problems are stored. In general, given two solutions from consecutive time steps, we let  $d(S_i^{(t)}, S_i^{(t+1)})$  denote the (problem-specific) *recourse* incurred by the  $i$ 'th solution at time step  $t$  (see below for how to set  $d(\cdot, \cdot)$  for the problems we study). The *total recourse* of an algorithm is given by

$$\sum_t \sum_{i=1}^s d(S_i^{(t)}, S_i^{(t+1)}).$$

Next, we will present the concrete recourse lower bounds that we derive.

**Recourse Bounds for Knapsack.** In knapsack, the solutions  $S_i^{(t)}$  simply correspond to subsets of items which are contained in the knapsack. To measure the recourse, we set  $d(S_i^{(t)}, S_i^{(t+1)}) = |S_i^{(t)} \Delta S_i^{(t+1)}|$ , i.e., we consider the cardinality of the symmetric difference of the  $i$ 'th solution at time steps  $t$  and  $t + 1$ .

Our main result shows that for a fixed accuracy  $\epsilon$ , any dynamic  $(1 - \epsilon)$ -approximation algorithm must maintain  $\Omega(1/\epsilon)$  solutions or it must have recourse  $\Omega(\frac{n}{\epsilon})$ , even when only a single item is inserted.

► **Theorem 34.** *Let  $\epsilon \in (0, 1/2)$ . Assume  $s < \frac{1}{8\epsilon(1+2\epsilon)}$  and  $n \in \mathbb{N}$  is a sufficiently large multiple of  $s$ . Then any dynamic randomized  $(1 - \epsilon)$ -approximation algorithm for knapsack with  $s$  solutions must have recourse  $\Omega(\frac{n}{s})$ . This holds even for a single item insertion.*



**Recourse Bounds for  $k$ -Balanced Partitioning.** In  $k$ -balanced partitioning, each solution  $S_i^{(t)}$  consists of  $k$  clusters  $V_1^{(i,t)}, \dots, V_k^{(i,t)}$  that partition the set of vertices  $V$ . To measure the recourse, we set  $d(S_i^{(t)}, S_i^{(t+1)}) = \sum_{j=1}^k |V_j^{(i,t)} \Delta V_j^{(i,t+1)}|$ , i.e., we consider the total number of vertices that change their set  $V_j^{(i,\cdot)}$  from time  $t$  to  $t+1$ .

Our main result shows that for any  $C$  and fixed  $\epsilon$ , any algorithm that maintains a  $(C, 1 + \epsilon)$ -approximate solution must use  $\Omega(1/\epsilon)$  solutions or it must have *amortized* recourse  $\Omega(\epsilon^2 \frac{n}{k})$ , even when only  $O(1/\epsilon)$  edges are inserted. Here, the amortized recourse refers to the total recourse divided by the total number of update operations.

► **Theorem 35.** *Let  $C > 0$  be arbitrary and  $\epsilon \in (0, 1/2)$ . Assume  $k \geq 4$  and  $s < \frac{1}{4\epsilon}$ . Then any dynamic randomized  $(C, 1 + \epsilon)$ -approximation algorithm for  $k$ -balanced partitioning with  $s$  solutions must have amortized recourse  $\Omega(\epsilon^2 \frac{n}{k})$ . This holds even for  $O(1/\epsilon)$  edge insertions.*

## F.1 Proof of Theorem 34

We prove Theorem 34. We use Yao's principle [63], i.e., we consider a deterministic algorithm and give a distribution over inputs, showing that in expectation the algorithm will have recourse  $\Omega(\frac{n}{s})$ .

We consider an instance in which initially we have  $n$  items, and each item  $i$  has weight  $w_i = 1$  and price  $p_i = 1$ . We refer to these items as *small* items. We set the budget of our knapsack to  $B = n$ . Note that in this instance,  $\text{OPT} = n$  because all small items fit into the knapsack.

Now we sample an integer  $j$  uniformly at random from  $[2s - 1] = \{0, \dots, 2s - 1\}$ , and we set  $k = i \cdot \frac{n}{2s} + \frac{n}{4s}$ . We insert a single *heavy* item with  $p = n - k + 2\epsilon n$  and  $w = n - k$ . Note that after inserting the heavy item, we have that  $\text{OPT} = n + 2\epsilon n$  since the optimal solution consists of the heavy item and  $k$  small items.

We let  $S_1, \dots, S_s$  denote the solutions maintained by the algorithm *before* the heavy item was inserted and we let  $S'_1, \dots, S'_s$  denote the solutions *after* the heavy item was inserted. We write  $\text{small}(S_i)$  to denote the number of small items in solution  $S_i$ .

In the following, we will show that any  $(1 - \epsilon)$ -approximate solution  $S'_i$  must contain the heavy item and “almost”  $k$  small items. However, we will also show that with constant probability all solutions  $S_i$  had “much less” or “much more” than  $k$  small items initially. This then gives that obtaining any  $(1 - \epsilon)$ -approximate solution must incur high recourse.

We follow this proof strategy in reverse order. We start by showing that with constant probability, all  $S_i$  have “much less” or “much more” than  $k$  small items.

► **Lemma 36.** *With probability at least  $1/2$ , it holds that  $|k - \text{small}(S_i)| \geq \frac{n}{4s}$  for all  $i \in [s]$ .*

**Proof.** Suppose that we partition the set  $\{1, \dots, n\}$  into  $2s$  consecutive intervals, each of length  $\frac{n}{2s}$ . Note that  $k$  is the middle point of one of these intervals. Furthermore, as there are only  $s$  solutions and  $2s$  intervals, at least half of the intervals do not contain a number from the set  $\{\text{small}(S_i) : i = 1, \dots, s\}$ ; we call these intervals *empty*. Thus, with probability at least  $1/2$ ,  $k$  is the middle point of an empty interval. If the interval containing  $k$  is empty, then  $k$  has distance at least  $\frac{n}{4s}$  to  $\text{small}(S_i)$  for all  $i = 1, \dots, s$ . ◀

Next, recall that  $S'_1, \dots, S'_s$  are the solutions maintained by the algorithm after the insertion of the heavy item. We show that any  $(1 - \epsilon)$ -approximate solution must contain the heavy item and some small items.

► **Lemma 37.** *Suppose that  $S'_i$  is a  $(1 - \epsilon)$ -approximate solution. Then  $S'_i$  contains the heavy item and a positive number of small items.*

**Proof.** First, suppose that a solution  $S'_i$  only contains small items. Then its total price is at most  $n$ . However, we have that

$$(1 - \epsilon) \text{OPT} = (1 - \epsilon)(1 + 2\epsilon)n > n,$$

where we used that  $\epsilon < 1/2$ . Hence,  $S'_i$  is not a  $(1 - \epsilon)$ -approximate solution.

Second, suppose that  $S'_i$  only contains the heavy item. Then we have that

$$\begin{aligned} (1 - \epsilon) \text{OPT} &= (1 - \epsilon)(p + k) \\ &= p + k - \epsilon(p + k) \\ &\geq p + \frac{n}{4s} - \epsilon(1 + 2\epsilon)n \\ &> p + 2\epsilon(1 + 2\epsilon)n - \epsilon(1 + 2\epsilon)n \\ &> p, \end{aligned}$$

where we used that  $k \geq \frac{n}{4s}$ ,  $p + k = n + 2\epsilon n$  and  $s < \frac{1}{8\epsilon(1+2\epsilon)}$ . Therefore, a solution containing only the heavy item is not  $(1 - \epsilon)$ -approximate. ◀

Next, we show that any  $(1 - \epsilon)$ -approximate solution must contain “almost”  $k$  small items.

► **Lemma 38.** *Suppose  $S'_i$  is a  $(1 - \epsilon)$ -approximate solution. Then*

$$k - \frac{n}{8s} \leq \text{small}(S'_i) \leq k.$$

**Proof.** The upper bound follows from the fact  $S'_i$  must contain the heavy item (by Lemma 37) of weight  $n - k$  and then it can only include  $k$  small items since the budget constraint is set to  $B = n$ .

To prove the lower bound, note that since  $S'_i$  is a  $(1 - \epsilon)$ -approximate solution, we have that its solution has value

$$p + \text{small}(S'_i) \geq (1 - \epsilon) \text{OPT} = (1 - \epsilon)(p + k).$$

Hence, we get that

$$\begin{aligned} \text{small}(S'_i) &\geq k - \epsilon(p + k) \\ &= k - \epsilon(1 + 2\epsilon)n \\ &> k - \frac{n}{8s}, \end{aligned}$$

where we used that  $s < \frac{1}{8\epsilon(1+2\epsilon)}$ . ◀

To finish the proof of the theorem, we condition on the event from Lemma 36, i.e., we have that  $|k - \text{small}(S_i)| \geq \frac{n}{4s}$  for all  $i \in [s]$ .

Now consider any solution  $S'_i$  after the insertion of the heavy item. If  $S'_i$  is not  $(1 - \epsilon)$ -approximate, we can ignore  $S'_i$ . If  $S'_i$  is  $(1 - \epsilon)$ -approximate then it satisfies  $k - \frac{n}{8s} \leq \text{small}(S'_i) \leq k$  by Lemma 38. Since we are assuming the event from Lemma 36, the algorithm had to insert/delete at least  $\frac{n}{8s}$  small items into/from  $S_i$  to obtain  $S'_i$ .

Since the event from Lemma 36 occurs with probability at least  $1/2$  and the above argument holds for all  $(1 - \epsilon)$ -approximate solutions  $S'_i$ , we have that the expected recourse is  $\Omega(\frac{n}{s})$ .

## F.2 Proof of Theorem 35

We prove Theorem 35. Again, we apply Yao's principle [63], i.e., we consider a deterministic algorithm and give a distribution over inputs, showing that in expectation the algorithm will have amortized recourse  $\Omega(\epsilon^2 \frac{n}{k})$ .

We consider a graph with  $n$  vertices. Our initial instance consists of  $\frac{k}{2\epsilon}$  star graphs, each of which contains  $2\epsilon \frac{n}{k}$  vertices. Note that here an optimal solution places the vertices from exactly  $\frac{1}{2\epsilon}$  star graphs into each partition  $V_j$ ; there are no edges between vertices from different  $V_j$  and hence the optimal cut-value is zero. Hence, the solution of any  $(C, 1 + \epsilon)$ -approximate solution must also have cut-value zero.

In the update phase, we sample  $s$  edges between the central nodes of the star graphs uniformly at random and insert them into the graph. Note that after the insertion of the edges, we connected at most  $s$  star graphs and the largest connected component has size at most  $s \cdot 2\epsilon \frac{n}{k} \leq \frac{1}{2} \frac{n}{k}$ , where we used that  $s \leq \frac{1}{4\epsilon}$ . Hence, the optimal solution still has cut-value zero and thus any  $(C, 1 + \epsilon)$ -approximate solution must have cut-value zero.

Next, let us analyze the recourse of an algorithm which starts with initial solutions  $S_1^{(0)}, \dots, S_s^{(0)}$ . In a first step, we show that solutions which at time 0 splits one of the star graphs up "too much" must entail high recourse. In a second step, we consider all other solutions and show that our insertions still trigger high recourse in expectation.

We say that a solution  $S_i^{(0)} = \{V_1^{(i,0)}, \dots, V_k^{(i,0)}\}$  is *useful* if for all star graphs  $H$ , it holds that there exists an index  $j$  such that  $V_j^{(i,0)}$  contains at least  $\epsilon \frac{n}{k}$  vertices from  $H$  and at most  $\epsilon \frac{n}{k}$  vertices are placed in  $\bigcup_{j' \neq j} V_{j'}^{(i,0)}$ . Given a star graph  $H$  and a solution  $S_i^{(0)}$ , we write  $j(H, i)$  to denote the index  $j$  such that  $V_j^{(i,0)}$  contains at least  $\epsilon \frac{n}{k}$  vertices from  $H$ . If a solution is not useful, we call it *useless*.

First, consider solutions which are useless. There exist two cases. *Case A:* Suppose there exists a star graph  $H$  such that for all indices  $j$  it holds that  $\bigcup_{j' \neq j} V_{j'}^{(i,0)}$  contains more than  $\epsilon \frac{n}{k}$  vertices from  $H$ . Observe that if the algorithm wants to use this solution after the edge insertions finished, it must ensure that the cut-value is zero. Thus it must move at least  $\epsilon \frac{n}{k}$  vertices to one of the  $V_j^{(i,0)}$  which requires  $\left| \bigcup_{j' \neq j} V_{j'}^{(i,0)} \right| \geq \epsilon \frac{n}{k}$  vertex moves. *Case B:* Suppose there exists a star graph  $H$  such that for all  $j$  it holds that  $V_j^{(i,0)}$  contains less than  $\epsilon \frac{n}{k}$  vertices from  $H$ . Using that  $\epsilon \in (0, \frac{1}{2})$ , also in this case the algorithm must move at least  $(1 - \epsilon) \frac{n}{k} \geq \epsilon \frac{n}{k}$  vertices such that eventually all of  $H$  is contained in the same set  $V_j^{(i,s)}$  when the updates finished. We conclude that for useless solutions our theorem holds after amortizing over  $s \leq \frac{1}{\epsilon}$  insertions.

Second, for the remainder of the proof consider only solutions  $S_i^{(0)} = \{V_1^{(i,0)}, \dots, V_k^{(i,0)}\}$  which are useful. Observe that when we insert an edge between two star graphs  $H_1$  and  $H_2$ , then if  $j(H_1, i) \neq j(H_2, i)$  the algorithm must move at least  $\epsilon \frac{n}{k}$  vertices to ensure that after the  $s$  insertions finished, all vertices from  $H_1$  and  $H_2$  are placed in the same set  $V_j^{(i,s)}$  for some  $j$ . We call such an insertion *expensive* for solution  $i$ .

Observe that if our edge insertions are such that they contain an expensive insertion for all solutions, then updating any solution  $S_i^{(0)}$  such that  $S_i^{(s)}$  is  $(C, 1 + \epsilon)$ -approximate will incur recourse at least  $\epsilon \frac{n}{k}$ . The rest of our proof is devoted to showing that with constant probability this event occurs. This will prove the theorem.

We start by considering a fixed solution  $S_i^{(0)}$  and a single random edge insertion between randomly picked star graphs  $H_1$  and  $H_2$ . Recall that there are  $\frac{k}{2\epsilon}$  star graphs in total. Furthermore, we have that  $\left| V_j^{(i,0)} \right| \leq (1 + \epsilon) \frac{n}{k}$  for all  $j$  and thus for each  $j$  there can be at most  $\frac{(1+\epsilon)}{\epsilon}$  star graphs  $H$  with  $j = j(H, i)$ . Hence, for the probability that the edge insertion

is expensive we get that

$$\begin{aligned}
\Pr(j(H_1, i) \neq j(H_2, i)) &= 1 - \Pr(j(H_1, i) = j(H_2, i)) \\
&\geq 1 - \frac{(1 + \epsilon)/\epsilon}{k/(2\epsilon)} \\
&= 1 - \frac{2(1 + \epsilon)}{k} \\
&\geq \frac{1}{2},
\end{aligned}$$

where we used that  $k \geq 4$ .

Next, we consider a fixed solution  $S_i^{(0)}$  and  $s$  edge insertions between star graphs which were picked independently and uniformly at random. Then with probability at least  $1 - 2^{-s}$ , at least one of these edge insertions is expensive for solution  $i$ .

Finally, observe that probability that for *all* solutions  $i$  there exists an expensive edge insertion is at least

$$\begin{aligned}
(1 - 2^{-s})^s &= \exp(s \ln(1 - 2^{-s})) \\
&\geq 1 + s \ln(1 - 2^{-s}) \\
&\geq 1 - s2^{-s} \\
&\geq \frac{1}{4},
\end{aligned}$$

where we used that  $\exp(x) \geq 1 + x$  for all  $x \in \mathbb{R}$ , the Taylor expansion of  $\ln(x)$  for  $x$  close to 1 and the fact that  $s2^{-s} \leq \frac{3}{4}$  for all  $s$ .

We conclude that with constant probability, for all solutions  $i$  there exists an expensive edge insertion. In this case, the algorithm has total recourse at least  $\epsilon \frac{n}{k}$ . Hence, the expected total recourse of the algorithm is  $\Omega(\epsilon \frac{n}{k})$ . Since we only performed  $s$  edge insertions, this gives an *amortized* recourse of  $\Omega(\epsilon^2 \frac{n}{k})$ .

## **G** Non-Monotone Functions and $\ell_\infty$ -Necklace Alignment

So far we have only considered *monotone* piecewise constant functions. Now we will generalize some of our results to *piecewise constant functions with multiple non-monotonicities* and provide the details in Section G.1. We also derive new approximation algorithms for the  $\ell_\infty$ -necklace problem in Section G.2. In particular, for  $\ell_\infty$ -necklace we present the first approximation algorithm *with near-linear running time* with additive error  $\epsilon$ . We also present the first *dynamic* approximation algorithm for this problem which achieves additive error  $\epsilon$  and has update time  $O((1/\epsilon)^2 \log(1/\epsilon))$ ; the algorithm has preprocessing time  $O(1)$  when starting with empty vectors  $x$  and  $y$  and requires *sublinear space*  $O(1/\epsilon)$ . See Theorem 44 for the details of our results.

### **G.1** Piecewise Constant Functions With Non-Monotonicities

We now show that we can perform efficient operations on piecewise constant functions even when these functions contain non-monotonicities. However, the running times of our subprocedures will typically have some dependency on the number of non-monotonicities of the function.

Let us formalize our notion of non-monotonicities. We say that a function  $f: [0, t] \rightarrow [0, W] \cup \{-\infty, \infty\}$  has  $k$  *monotone segments* if there exist values  $0 = x_0 < x_1 < \dots < x_k = t$

such that on each interval  $[x_i, x_{i+1})$ ,  $f$  is monotone. Here, we require that either  $f$  is monotonically decreasing on *all* segments or it is monotonically increasing on *all* segments. Note that a monotone function has  $k = 1$  monotone segments (by setting  $x_0 = 0$  and  $x_1 = t$ ) and that the points  $x_1, \dots, x_{k-1}$  can be viewed as the points in which  $f$  is non-monotone.

One crucial operations will again be rounding. However, unlike previously we will mostly talk about rounding to *multiples* of  $\delta$  instead of rounding to powers of  $1 + \delta$ . This will be convenient for our applications to  $\ell_\infty$ -necklace later. We will also briefly mention how to extend our results from this subsection to the setting in which we round to powers of  $1 + \delta$ .

Next, let  $\delta > 0$  and consider a simple rounding function that rounds down to multiples of  $\delta$ . More concretely, for  $y \in \mathbb{R}$  we set  $\lfloor y \rfloor_\delta^* = \max\{i \cdot \delta : i \cdot \delta \leq y, i \in \mathbb{Z}\}$  and we follow the convention that  $\lfloor -\infty \rfloor_\delta^* = -\infty$  and  $\lfloor \infty \rfloor_\delta^* = \infty$ . We also extend the rounding operation to functions  $f: [0, t) \rightarrow [0, W] \cup \{-\infty, \infty\}$  by defining  $\lfloor f \rfloor_\delta^*: [0, t) \rightarrow [0, W] \cup \{-\infty, \infty\}$  to be the function with  $\lfloor f \rfloor_\delta^*(x) = \lfloor f(x) \rfloor_\delta^*$  for all  $x \in [0, t)$ . Next, we show that the function  $\lfloor f \rfloor_\delta^*$  can be computed efficiently and that it has only few pieces.

► **Lemma 39.** *Let  $\delta > 0$  and let  $f: [0, t) \rightarrow [0, W] \cup \{-\infty, \infty\}$  be a piecewise constant function with  $p$  pieces and  $k$  monotone segments. Then we can compute the function  $\lfloor f \rfloor_\delta^*$  in time  $O(p \log p)$  and  $\lfloor f \rfloor_\delta^*$  has  $O(k \cdot W/\delta)$  pieces.*

**Proof.** Let  $(x_1, y_1), \dots, (x_p, y_p)$  denote the list representation of  $f$ . We construct the list representation  $(x'_1, y'_1), \dots, (x'_p, y'_p)$  of  $\lfloor f \rfloor_\delta^*$ . For all  $i = 1, \dots, p$ , we set  $x'_i = x_i$  and  $y'_i = \lfloor y_i \rfloor_\delta^*$ . After that, we merge all consecutive pieces that have the same  $y'_i$ -values; this can be done exactly as in the pruning step described in the proof of Lemma 6. Since  $f$  takes values in  $[0, W] \cup \{-\infty, \infty\}$ , there are  $O(W/\delta)$  choices for multiples of  $\delta$  in  $[0, W]$ . In particular, on each monotone segment of  $f$ ,  $\lfloor f \rfloor_\delta^*$  has  $O(W/\delta)$  pieces. Since  $f$  has  $k$  monotone segments this implies that  $\lfloor f \rfloor_\delta^*$  has  $O(k \cdot W/\delta)$  pieces in total. Note that all operations from above can be performed in linear time and the running time bound stems from the fact that we also need to store the pieces in a binary search tree. ◀

Next, we show that we can compute the  $(\min, +)$ -convolution of two piecewise constant functions in time that is quadratic in the number of their pieces. The lemma generalizes the result from Lemma 7 because we drop the assumption that one of the functions needs to be monotone (but this comes at the cost of a more complicated proof). We prove the lemma in Section G.1.1.

► **Lemma 40.** *Let  $f_1, f_2: [0, t) \rightarrow [0, W] \cup \{-\infty, \infty\}$  be piecewise constant functions which have at most  $p$  pieces. Then we can compute  $f = f_1 \oplus f_2$  in time  $O(p^2 \log p)$  and  $f$  has  $O(p^2)$  pieces.*

By combining the two lemmas above, we can show that we can efficiently compute additive approximations of  $(\min, +)$ -convolutions even in the case of non-monotonocities. More concretely, we say that  $f: [0, t) \rightarrow [0, W] \cup \{-\infty, \infty\}$  is an *additive  $\epsilon$ -approximation* of  $g: [0, t) \rightarrow [0, W] \cup \{-\infty, \infty\}$  if  $g(x) - \epsilon \leq f(x) \leq g(x)$  for all  $x \in [0, t)$ . Now we obtain the following theorem.

► **Theorem 41.** *Let  $f, g: [0, t) \rightarrow [0, W] \cup \{-\infty, \infty\}$  be two functions with  $k$  monotone segments and suppose we have already computed  $\lfloor f \rfloor_\delta^*$  and  $\lfloor g \rfloor_\delta^*$ . Then the function  $(\lfloor f \rfloor_\delta^*) \oplus (\lfloor g \rfloor_\delta^*)$  is an additive  $2\delta$ -approximation of  $f \oplus g$ , has at most  $O((k \cdot W/\delta)^2)$  pieces and can be computed in time  $O((k \cdot W/\delta)^2 \log((k \cdot W/\delta)^2))$ .*

**Proof.** The approximation ratio follows from the triangle inequality. The claims about the number of pieces and the running time follow from combining Lemma 39 and Lemma 40. ◀

We note that by stating Lemma 39 for the rounding operation  $\lceil \cdot \rceil_{1+\delta}$  that rounds to powers of  $1 + \delta$  (see Lemma 6), we can obtain the following version of Theorem 41.

► **Theorem 42.** *Let  $f, g: [0, t] \rightarrow [0, W] \cup \{-\infty, \infty\}$  be two functions with  $k$  monotone segments and suppose we have already computed  $\lceil f \rceil_{1+\delta}$  and  $\lceil g \rceil_{1+\delta}$ . Then  $(\lceil f \rceil_{1+\delta}) \oplus (\lceil g \rceil_{1+\delta})$  is a  $(1+\delta)$ -approximation of  $f \oplus g$ , has at most  $O((k \cdot \log_{1+\delta}(W))^2)$  pieces and can be computed in time  $O((k \cdot \log_{1+\delta}(W))^2 \log((k \cdot \log_{1+\delta}(W))^2))$ .*

This result generalizes our previous method of first rounding a monotone function via Lemma 6 and then applying the efficient convolution from Lemma 7. More concretely, observe that monotone functions have one monotone segment and, thus, after rounding both functions, our algorithm from Lemma 7 computes the  $(\min, +)$ -convolution in time  $O(\log_{1+\delta}^2(W) \log \log_{1+\delta}(W))$  which is the same running time that we obtain by combining the two lemmas above. Hence, the algorithm from Theorem 42 matches this result for  $k = 1$  and it generalizes it when we apply it for  $k > 1$ .

### G.1.1 Proof of Lemma 40

We assume that  $f_i$  for  $i = 1, 2$  is given as a doubly linked list  $(x_1^i, y_1^i), \dots, (x_p^i, y_p^i)$  such that  $x_j^i < x_{j+1}^i$  for all  $1 \leq j < p$ . We will output  $f$  in the same representation.

To compute  $f$  we will make use of the following *non-overlapping interval data structure (NOI)*. Let  $[a, b]$  and  $[a', b']$  be two subsets of the real line. We call each of them an *interval* and say that they *overlap* if  $[a, b] \cap [a', b'] \neq \emptyset$ . We say that an interval  $[a, b]$  is *empty* if  $a \geq b$ . The NOI data structure stores a set  $S$  of non-overlapping, non-empty intervals  $I = [a, b]$  and supports the following operations:

- *ClosestLargerInterval*( $z$ ), which given a number  $z$  returns the interval  $[a, b]$  together with a Boolean value *bool*. If *bool* is true, then  $z \leq b$  and there is no interval  $[a', b']$  in  $S$  with  $z \leq b' < b$ . Note that it is possible that  $z$  belongs to  $[a, b]$ . If *bool* is false, then there exists no interval  $[a, b]$  with  $z \leq b$  and the returned values for  $a$  and  $b$  are undefined.
- *InsertInterval*( $a, b$ ), which inserts the interval  $[a, b]$  into  $S$ , merging it with any interval that it overlaps with and updating  $S$  accordingly.

There exists an efficient implementation of such a data structure as stated in the next claim, which we prove at the end of this section.

▷ **Claim 43.** There exists an implementation of the non-overlapping interval data structure such that any sequence of  $q$  operations takes time  $O(q \log q)$ .

We compute  $f$  as follows. Note that the function values of  $f_1$  and of  $f_2$  are constant over each 2-dimensional rectangle whose corners are  $(x_s^1, x_t^2)$ ,  $(x_s^1, x_{t+1}^2)$ ,  $(x_{s+1}^1, x_t^2)$ , and  $(x_{s+1}^1, x_{t+1}^2)$  for any  $1 \leq s \leq p$  and  $1 \leq t \leq p$ . We call this rectangle  $R_{st}$  and denote by  $[x_s^1 + x_t^2, x_{s+1}^1 + x_{t+1}^2]$  the *range* of the rectangle  $R_{st}$  and by  $y_s^1 + y_t^2$  the *function value* of the rectangle, where we assume that  $\infty + y$  with  $y \in W_\infty$  equals  $\infty$ . There are  $K^2$  such rectangles.

Now note that for any value  $x$  with  $x_s^1 + x_t^2 \leq x \leq x_{s+1}^1 + x_{t+1}^2$ , i.e.,  $x$  is in the range of the rectangle  $R_{st}$ , the function value  $y_s^1 + y_t^2$  is one of the sums that occurs in the computation of  $f(x) = \min_{\bar{x}} \{f_1(\bar{x}) + f_2(x - \bar{x})\}$ . We will compute  $f(x)$  (for all values  $x$  “simultaneously”) by comparing the function values of all rectangles  $R_{st}$  to whose range  $x$  belongs. The main observation that we exploit is the following: *As we will consider the rectangles by decreasing function values, the first rectangle (in this order) to whose range a value  $x$  belongs is the rectangle whose function value equals  $f(x)$ .*

Thus, when processing a rectangle, we need to determine all ranges, i.e., subintervals of  $[0, t]$ , to which no function value has yet been assigned. To do so, we use the NOI data structure to store the intervals of all values  $x$  for which we have already assigned a function value. Furthermore we use a balanced binary search tree  $\mathcal{B}$  that stores at its leaves every interval to which a function value has already been assigned, together with its (constant) function value. Specifically, we will store these ranges in the leaves of  $\mathcal{B}$ , ordered by their smaller boundary value  $x'$ . The difference between the two is that the NOI data structures merges overlapping intervals, no matter what their function value is, while every interval stored as a leaf of  $\mathcal{B}$  has the *same* function value, i.e., has a constant  $f$ -value.

To be precise we proceed as follows: We first generate all rectangles  $R_{st}$  by iterating over the lists of  $f_1$  and  $f_2$  and sort them by non-decreasing order of their function value. This takes time  $O(p^2 \log p)$ . Then we process the rectangles in this order. To do so, we first initialize an empty NOI data structure as well as an empty balanced binary search tree  $\mathcal{B}$ . Next we describe how to process the rectangles. Let  $R_{st}$  be the next rectangle to be processed. We execute the following steps for  $R_{st}$ :

1.  $z = x_s^1 + x_t^2$
2.  $(a, b, \text{bool}) = \text{ClosestLargerInterval}(z)$
3. **while**  $\text{bool}$  is true and  $b < x_{s+1}^1 + x_{t+1}^2$  **do**
  - a. if  $z \notin [a, b]$  then insert the interval  $[z, a]$  together with the function value of  $R_{st}$  into  $\mathcal{B}$ .
  - b.  $z = b$
  - c.  $(a, b, \text{bool}) = \text{ClosestLargerInterval}(z)$
4. If  $\text{bool}$  is true then insert the interval  $[z, a]$  together with the function value of  $R_{st}$  into  $\mathcal{B}$ ; else insert the interval  $[z, x_{s+1}^1 + x_{t+1}^2]$  together with the function value of  $R_{st}$  into  $\mathcal{B}$ .
5.  $\text{InsertInterval}(x_s^1 + x_t^2, x_{s+1}^1 + x_{t+1}^2)$ .

Once all rectangles have been processed, we traverse the leaves of  $\mathcal{B}$  in order and connect them by a doubly linked list to create an (ordered) list representation of the function  $f$ . As we process the rectangles in increasing order of function value this guarantees that for each value  $x$  the smallest function value of any rectangle  $R_{st}$  is returned as  $f(x)$ .

Note that each insertion into  $\mathcal{B}$  takes time  $O(\log p)$  and the number of calls to the NOI data structure is proportional to the number of rectangles plus the number of intervals merged in the NOI data structure. As processing a rectangle creates at most one new interval, and merged intervals are never separated again, the number of interval merges is at most the number of rectangles. Thus, there are at most  $p$  interval merges and at most  $2p^2$  insertions into  $\mathcal{B}$ . Hence, the total running time for the above algorithm is  $O(p^2 \log p)$  plus the time for the NOI data structure, which, by Claim 43, is also  $O(p^2 \log p)$  as  $q = O(p^2)$ .

We still have to prove Claim 43.

**Proof of Claim 43.** We implement the NOI data structure with a balanced binary search tree. The leaves store the non-overlapping intervals, ordered by their upper endpoint.

The  $\text{ClosestLargerInterval}(z)$  operation searches for the interval  $[a, b]$  such that  $b$  is the smallest upper endpoint of an interval that is at least  $z$ . If no such interval exists,  $\text{bool}$  is set to false, otherwise it is set to true and  $[a, b]$  is returned as interval. Note that finding  $[a, b]$  takes time  $O(\log q)$ , as  $q$  is the maximum number of intervals stored in the balanced binary tree.

The  $\text{InsertInterval}(a, b)$  operation first executes a  $\text{ClosestLargerInterval}(a)$  operation. Let  $(a', b', \text{bool})$  be the result. If  $\text{bool}$  is false, then the interval  $[a, b]$  is inserted as new interval and the procedure terminates. Otherwise the interval  $[a', b']$  is the interval with smallest upper endpoint such that  $a \leq b'$ . Note that  $[a', b']$  might overlap with  $[a, b]$  and we test for this next. If  $b < a'$  then a leaf with range  $[a, b]$  is inserted into the balanced search tree and

$InsertInterval(a, b)$  terminates. Otherwise ( $b \geq a'$ ), let  $L$  be the leaf of the balanced search tree that stores  $[a', b']$ . If  $b \leq b'$ , the two intervals are merged by updating  $L$  to store the interval  $[\min(a, a'), b']$  and  $InsertInterval(a, b)$  terminates. If, however,  $b > b'$ , it is possible that the new interval  $[a, b]$  overlaps with even more intervals in  $S$ . Thus, we execute the following steps:

1.  $z = b'$
2.  $(a'', b'', bool) = ClosestLargerInterval(z)$
3. **while**  $bool$  is true **do**
  - a. If  $b < a''$  then the leaf  $L$  is updated to store the interval  $[\min(a, a'), b]$  and  $InsertInterval$  terminates.
  - b. The leaf storing the interval  $[a'', b'']$  is removed from the balanced search tree.
  - c. If  $b \leq b''$  then the leaf  $L$  is updated to store the interval  $[\min(a, a'), b'']$  and  $InsertInterval$  terminates.
  - d. Otherwise,  $z = b''$  and  $(a'', b'', bool) = ClosestLargerInterval(z)$ .
4.  $L$  is updated to store the interval  $[\min(a, a'), b]$ .

Note that this algorithm merges all intervals that overlap with  $[a, b]$  into one interval and updates the balanced search tree accordingly.

Let  $t$  be the number of iterations executed by  $InsertInterval(x, y)$ . The running time is  $O((t + 1) \log q)$  as each iteration executes one call to  $ClosestLargerInterval$ , one deletion of a leaf in the balanced binary tree, and at most one modification of a label at a leaf. Every such iteration decreases the number of leaves in the balanced binary tree by 1. Furthermore, each call to  $InsertInterval$  that does not execute any iterations of the above while-loop increases the number of leaves by at most 1 and there is no other operation that modifies the number of leaves. As there are at most  $q$  calls to  $InsertInterval$ , the while-loop can be executed at most  $q$  times *over all calls to  $InsertInterval$* , each taking time  $O(\log q)$ . Thus, the total running time for  $q$  calls to  $InsertInterval$  is  $O(q \log q)$ . ◀

## G.2 $\ell_\infty$ -Necklace Alignment

Using our techniques from above, we present a novel approximation algorithm for the  $\ell_\infty$ -necklace alignment problem [14, 58]. In this problem, the input consists of two necklaces represented as two *sorted* vectors of  $n$  real numbers,  $x = \langle x_0, x_1, \dots, x_{n-1} \rangle$  and  $y = \langle y_0, y_1, \dots, y_{n-1} \rangle$ , where the  $x_i, y_i \in [0, 1)$  represent points on the unit-circumference circle. We will sometimes refer to the elements  $x_i$  and  $y_j$  as *beads*.

We define the distance between two beads  $x_i$  and  $y_j$  by the minimum of the clockwise and counterclockwise distances along the circumference of the unit-perimeter circular necklaces, i.e., we set

$$d^\circ(x_i, y_j) = \min\{|x_i - y_j|, 1 - |x_i - y_j|\}.$$

In the  $\ell_\infty$ -necklace alignment problem, we need to find an *offset*  $c \in [0, 1)$  and a *shift*  $s \in [n + 1]$  that minimize

$$\max_{i=0}^{n-1} (d^\circ((x_i + c) \bmod 1, y_{(i+s) \bmod n})).$$

In the above definition, the offset  $c$  encodes how much we rotate the first necklace clockwise relative to the second necklace. Additionally, the shift  $s$  defines a perfect matching between the beads such that bead  $i$  of the first necklace is matched with bead  $(i + s) \bmod n$  of the second necklace.



Bremner et al. [14] showed that the  $\ell_\infty$ -necklace alignment problem can be solved *exactly* in time  $\tilde{O}(n^2)$ . We complement this by showing that we can compute a solution with additive error  $\epsilon$  in time  $\tilde{O}(n + \epsilon^{-2})$ .

We also consider the dynamic version of the problem in which beads are inserted and deleted. More concretely, we assume that initially  $x$  and  $y$  are empty and we offer the following update operations:

- *Insert*( $i, \alpha, \beta$ ) which inserts  $\alpha \in [0, 1)$  into  $x$  at the  $i$ 'th position and it further inserts  $\beta \in [0, 1)$  into  $y$  at the  $i$ 'th position. We require that after the insertion,  $x$  and  $y$  are still ordered.
- *Delete*( $i$ ) which deletes  $x_i$  from  $x$  and  $y_i$  from  $y$ .

Note that both of these operations change the number of entries in  $x$  and  $y$  but they ensure that  $x$  and  $y$  always have the same length. We show that we can maintain a solution with additive error  $\epsilon$  using update time  $O(1/\epsilon^2 \log(1/\epsilon))$ . The preprocessing time is  $O(1)$  and the space usage is only  $O(1/\epsilon)$  which is sublinear in the size of the vectors  $x$  and  $y$ .

► **Theorem 44.** *Let  $\epsilon > 0$ . There exists a static algorithm for the  $\ell_\infty$ -necklace alignment problem that computes a solution with additive error  $\epsilon$  in time  $O(n + (1/\epsilon)^2 \log(1/\epsilon))$ . Furthermore, there exists a fully dynamic algorithm for the  $\ell_\infty$ -necklace alignment problem that maintains a solution with additive error  $\epsilon$  with update time  $O(1/\epsilon^2 \log(1/\epsilon))$  and preprocessing time  $O(1)$ ; the space usage of the algorithm is  $O(1/\epsilon)$ .*

To obtain the result for the dynamic algorithm, we show that for vectors  $A, B \in \mathbb{R}^n$  that are undergoing element insertions and deletions, we can dynamically maintain an approximation of the  $(\min, +)$ -convolution  $A \oplus B$ . We expect that this result will have further applications. The proof of the theorem follows from Propositions 45 and 48 below.

## G.2.1 The Static Algorithm

Now we consider our static algorithm and prove the following proposition.

► **Proposition 45.** *There exists a static algorithm for the  $\ell_\infty$ -necklace alignment problem that computes a solution with additive error  $\epsilon$  in time  $O(n + (1/\epsilon)^2 \log(1/\epsilon))$ .*

We devote the rest of this subsection to the proof of the proposition.

**The Algorithm.** Our algorithm is rather simple and (up to the part in which we perform the rounding) it is the same as the one used by Bremner et al. [14]. Consider the input  $\epsilon$  (as error parameter),  $x = \langle x_0, x_1, \dots, x_{n-1} \rangle$  and  $y = \langle y_0, y_1, \dots, y_{n-1} \rangle$ . Now we set  $\delta = \epsilon/2$  and perform a single pass over  $x$  and  $y$  and apply the rounding function  $\lfloor \cdot \rfloor_\delta^*$  to each of the entries. While doing so, we compute the list representations of  $x$  and  $y$  (where we interpret  $x$  and  $y$  as functions from  $[0, n)$  to  $[0, 1)$ ) which have at most  $O(1/\delta)$  pieces (by applying Lemma 39 with  $W = 1$ ). Then we compute the vectors

$$\begin{aligned} x' &= \langle x_0, x_1, \dots, x_{n-1}, \underbrace{\infty, \dots, \infty}_{n \text{ times}} \rangle, \\ x'' &= \langle x_0, x_1, \dots, x_{n-1}, \underbrace{-\infty, \dots, -\infty}_{n \text{ times}} \rangle, \\ y' &= \langle y_{n-1}, y_{n-2}, \dots, y_0, y_{n-1}, y_{n-2}, \dots, y_0 \rangle, \end{aligned}$$

but we do not store them explicitly. Instead, we only store their list representations. We note that  $x'$  is a monotonically increasing vector,  $x''$  has two monotonically increasing segments and  $y'$  has two monotonically decreasing segments.

Next, we set  $a$  to the  $(\min, -)$ -convolution of  $x'$  and  $y'$  and we set  $b$  to the  $(\max, -)$ -convolution of  $x''$  and  $y'$  (we show below in Lemma 46 that we can compute these functions efficiently).

Finally, we set  $v = \frac{1}{2}(b - a)$  and return  $\min\{v_s : s \in [n]\}$  as the solution for our problem. We note that  $v$  can be efficiently computed via the list representations of  $a$  and  $b$  and we can also quickly find the minimum over the  $v_s$  by iterating over the list representation of  $v$ .

**Analysis.** Now we turn to the analysis of the algorithm above. We adapt the proof of Theorem 6 in Bremner et al. [14] for approximate solutions and argue how to implement it using piecewise constant functions.

We start by showing that we can compute  $(\min, -)$ -convolution and  $(\max, -)$ -convolution as efficiently as the classic  $(\min, +)$ -convolution.

► **Lemma 46.** *Let  $f$  and  $g$  be two piecewise constant functions with  $p$  pieces and suppose that  $g$  has  $k$  monotonically decreasing segments. Suppose that we can compute the  $(\min, +)$ -convolution of  $f'$  and  $g'$  in time  $t(p, k)$  if  $f'$  and  $g'$  have  $k$  monotonically decreasing segments. Then in time  $O(t(p, k) + p \log p)$  we can compute:*

- The  $(\max, -)$ -convolution of  $f$  and  $g$  if  $f$  has  $k$  monotonically increasing segments.
- The  $(\min, -)$ -convolution of  $f$  and  $g$  if  $f$  has  $k$  monotonically increasing segments.

**Proof.** First, suppose that we wish to compute the  $(\max, -)$ -convolution of two functions  $f$  and  $g$ . We show that we can compute the  $(\max, -)$ -convolution of  $f$  and  $g$  via the  $(\min, +)$ -convolution of  $-f$  and  $g$ . Indeed, for all  $x$  it holds that:

$$\begin{aligned} \max_{\bar{x} \in [0, x]} \{f(\bar{x}) - g(x - \bar{x})\} &= \max_{\bar{x} \in [0, x]} \{-(-f(\bar{x}) + g(x - \bar{x}))\} \\ &= - \min_{\bar{x} \in [0, x]} \{-f(\bar{x}) + g(x - \bar{x})\} \\ &= -(((-f) \oplus g)(x)). \end{aligned}$$

To see that the running time is correct, note that we can compute the list representation of  $-f$  in time  $O(p)$  and it takes  $O(p \log p)$  to update the binary search tree in which we store the pieces of  $-f$ . Furthermore,  $-f$  has  $k$  monotonically decreasing segments since  $f$  has  $k$  monotonically increasing segments. Thus, we can apply the efficient algorithm for  $(\min, +)$ -convolution in time  $t(p, k)$  on  $-f$  and  $g$ .

We can prove the result for  $(\min, -)$ -convolution similarly by computing a  $(\min, +)$ -convolution of  $g$  and  $-f$ . More concretely, for all  $x$  it holds that

$$\min_{\bar{x} \in [0, x]} \{f(\bar{x}) - g(x - \bar{x})\} = \min_{\bar{x} \in [0, x]} \{-f(\bar{x}) + g(x - \bar{x})\} = ((-f) \oplus g)(x),$$

where in the first step we used the symmetry of  $(\min, -)$ -convolution. The running time analysis is exactly as above. ◀

In the proof of Proposition 45 we need the following lemma. We will use the lemma to find the optimal offset  $c$  for a given shift  $s$ .

► **Lemma 47 (Fact 5 in [14]).** *Let  $z = \langle z_0, z_1, \dots, z_{n-1} \rangle$ . Then*

$$\min_{c \in \mathbb{R}} \max_{i=0}^{n-1} |z_i + c| = \frac{1}{2} \left( \max_{i=0}^{n-1} z_i - \min_{i=0}^{n-1} z_i \right)$$

and the minimizer for this quantity is given by  $c = -\frac{1}{2}(\min_{i=0}^{n-1} z_i + \max_{i=0}^{n-1} z_i)$ .

Next, we can prove Theorem 44.

**Proof of Theorem 44.** We prove the theorem in three steps. In Step 1, we will prove that we compute the correct result in the exact case (i.e., when we perform no rounding). This first part is essentially the same proof as in in Bremner et al. [14] but with more details. In Step 2, we argue about approximation guarantee of our algorithm. In Step 3, we prove its running time.

*Step 1: The Exact Case.* First, we use Theorem 2 of Bremner et al. [14] which states that if

$$\tilde{y} = \langle y_0, y_1, \dots, y_{n-1}, y_0, y_1, \dots, y_{n-1} \rangle$$

then

$$\min_{c,s} \max_{i=0}^{n-1} d^\circ((x_i + c) \bmod 1, y_{(i+s) \bmod n}) = \min_{c,s} \max_{i=0}^{n-1} d^-(x_i + c, \tilde{y}_{i+s}),$$

where  $d^-(a, b) = |a - b|$  for all  $a, b \in \mathbb{R}$ . Thus, instead of directly optimizing the original objective function  $\min_{c,s} \max_{i=0}^{n-1} d^\circ((x_i + c) \bmod 1, y_{(i+s) \bmod n})$ , we will consider the more convenient objective function  $\min_{c,s} \max_{i=0}^{n-1} d^-(x_i + c, \tilde{y}_{i+s})$  which involves no modulo operations.

Indeed, consider the new objective function and for all  $s \in [n]$  we define the vector  $z(s) \in \mathbb{R}^n$  such that  $z(s)_i = x_i - y_{(i+s) \bmod n}$ . Now we obtain that for the new objective function it holds that:

$$\begin{aligned} \min_{c,s} \max_{i=0}^{n-1} d^-(x_i + c, \tilde{y}_{i+s}) &= \min_{c,s} \max_{i=0}^{n-1} |x_i + c - y_{(i+s) \bmod n}| \\ &= \min_s \min_c \max_i |z(s)_i + c| \\ &= \min_s \frac{1}{2} \left( \max_i \{z(s)_i\} - \min_i \{z(s)_i\} \right), \end{aligned}$$

where in the first step we used the definition of  $d^-(\cdot, \cdot)$  and that  $\tilde{y}_k = y_{k \bmod n}$  for all  $k \in [2n]$ , in the second step we substituted the definition of  $z(s)_i$  and in the third step we applied Lemma 47.

The above implies that we need to compute the quantities  $\max_i \{z(s)_i\}$  and  $\min_i \{z(s)_i\}$  efficiently. Even more, consider the vector  $v \in \mathbb{R}^n$  with entries  $v_s = \frac{1}{2} (\max_i \{z(s)_i\} - \min_i \{z(s)_i\})$  and observe that the calculation above shows that the optimal objective function value is the same as the smallest entry in  $v$ . Therefore, in the following we show that we can compute  $v$  efficiently using the vectors  $a$  and  $b$  that we computed in our algorithm.

Recall the definitions of the two vectors  $x'$  and  $y'$ :

$$x' = \langle x_0, x_1, \dots, x_{n-1}, \underbrace{\infty, \dots, \infty}_{n \text{ times}} \rangle,$$

$$y' = \langle y_{n-1}, y_{n-2}, \dots, y_0, y_{n-1}, y_{n-2}, \dots, y_0 \rangle.$$

Now we let  $a \in \mathbb{R}^{2n}$  be the vector resulting from the  $(\min, -)$ -convolution of  $x'$  and  $y'$ , i.e.,  $a_k = \min_i \{x'_i - y'_{k-i}\}$  for all  $k \in [2n]$ . Now we observe that for each entry  $a_{n+s'}$  with  $s' \in [n]$ , it holds that

$$a_{n+s'} = \min_{i=0}^{n+s'} \{x'_i - y'_{n+s'-i}\} = \min_{i=0}^{n-1} \{x_i - y_{(i-s'-1) \bmod n}\},$$

where in the second step we used that  $x'_i = \infty$  for  $i \geq n$  and that  $y'_{n+s'-i} = y_{((n-1)-(n+s'-i)) \bmod n} = y_{(i-s'-1) \bmod n}$  since in  $y'$  we concatenated the entries of  $y$  twice but in reverse order. Now observe that if  $s' = n - 1 - s$  then

$$a_{2n-s-1} = a_{n+s'} = \min_{i=0}^{n-1} \{x_i - y_{(i-s'-1) \bmod n}\} = \min_{i=0}^{n-1} \{x_i - y_{(i+s) \bmod n}\} = \min_{i=0}^{n-1} \{z(s)_i\}.$$

Next, we define the vector  $x''$  such that:

$$x'' = \langle x_0, x_1, \dots, x_{n-1}, \underbrace{-\infty, \dots, -\infty}_{n \text{ times}} \rangle.$$

We let  $b$  denote the vector resulting from the  $(\max, -)$ -convolution of  $x''$  and  $y'$ , i.e.,  $b_k = \max_i \{x''_i - y'_{k-i}\}$  for all  $k \in [2n]$ . Now a similar argument as above shows that  $b_{2n-s-1} = \max_{i=0}^{n-1} \{z(s)_i\}$  for all  $s \in [n]$ . More concretely, for each entry  $b_{n+s'}$  with  $s' \in [n]$  it holds that

$$b_{n+s'} = \max_{i=0}^{n+s'} \{x'_i - y'_{n+s'-i}\} = \max_{i=0}^{n-1} \{x_i - y_{(i-s'-1) \bmod n}\},$$

where we used that  $x''_i = -\infty$  for  $i \geq n$  and the same argument relating the entries of  $y'$  and  $y$  as above. Thus, if  $s' = n - 1 - s$  then

$$b_{2n-s-1} = b_{n+s'} = \max_{i=0}^{n-1} \{x_i - y_{(i-s'-1) \bmod n}\} = \max_{i=0}^{n-1} \{x_i - y_{(i+s) \bmod n}\} = \max_{i=0}^{n-1} \{z(s)_i\}.$$

Combining the results above we get that  $v_s = \frac{1}{2}(b_{2n-s-1} - a_{2n-s-1})$  for all  $s \in [n]$ . Therefore, we get that the optimal objective function value is given by  $\min_s v_s = \min_s \frac{1}{2}(b_{2n-s-1} - a_{2n-s-1})$ . In other words, to compute the optimal objective function value it suffices to compute the difference  $\frac{1}{2}(b - a)$  and then to return the smallest entry in  $v$  with index between  $n$  and  $2n - 1$ .

*Step 2: Approximation Guarantees.* We argue that the algorithm returns an additive  $\epsilon$ -approximation. First, observe that in the algorithm all computations are performed exactly except for the rounding at the beginning. In the rounding process, we decrease each entry by at most  $\delta = \epsilon/2$ . Therefore, the triangle inequality implies that when we match bead  $x_i$  to bead  $y_{i+s}$ , the error that was introduced by the approximation is at most  $2\delta = \epsilon$ . Since in the objective function we are only interested in the maximum error over all matched beads, this implies that we obtain an additive  $\epsilon$ -approximation.

*Step 3: Running Time Analysis.* It is left to analyze the running time of our algorithm. Iterating over the input vectors  $x$  and  $y$ , rounding the entries and computing the list representation of  $x$  and  $y$  can be done in time  $O(n)$ . Recall that  $x$  and  $y$  have  $O(1/\delta)$  pieces. Therefore, we can also compute the vectors  $x'$ ,  $x''$  and  $y$  in time  $O(1/\delta \log(1/\delta))$ . Then Lemmas 46 and 40 imply that we can compute the  $(\min, -)$ -convolution and the  $(\min, +)$ -convolutions in time  $O(1/\delta^2 \log(1/\delta))$  and the resulting vectors have  $O(1/\delta^2)$  pieces. Finally, the vector  $v$  can be computed in time  $O(1/\delta^2 \log(1/\delta))$  and the minimum that we return can be found by simply iterating over the pieces of  $v$ . Since previously we have set  $\delta = \epsilon/2$ , this finishes the proof.  $\blacktriangleleft$

## G.2.2 The Dynamic Algorithm

We now give our extension to the dynamic setting of the  $\ell_\infty$ -necklace alignment problem in which there are insertions and deletions from  $x$  and  $y$ .

**► Proposition 48.** *Let  $\epsilon > 0$ . There exists a fully dynamic algorithm for the  $\ell_\infty$ -necklace alignment problem that maintains a solution with additive error  $\epsilon$  with update time  $O(1/\epsilon^2 \log(1/\epsilon))$  and preprocessing time  $O(1)$ ; the space usage of the algorithm is  $O(1/\epsilon)$ .*

**Proof.** In the preprocessing, we initialize  $x$  and  $y$  as empty vectors and store them as piecewise constant functions (as per Section 2) and we *do not* store them explicitly as vectors. Furthermore, we set  $\delta = \epsilon/2$ . These operations can be done in time  $O(1)$ .

Next, consider an operation  $Insert(i, \alpha, \beta)$  which asks to insert  $\alpha$  into  $x$  at the  $i$ 'th position and to insert  $\beta$  into  $y$  at the  $i$ 'th position. Since we are in the approximate setting, instead of inserting the exact values of  $\alpha$  and  $\beta$ , we insert  $\lfloor \alpha \rfloor_\delta^*$  into the  $i$ 'th position of  $x$  and  $\lfloor \beta \rfloor_\delta^*$  into the  $i$ 'th position of  $y$ . We perform these insertions by manipulating the list representations of  $x$  and  $y$ . We only describe how to perform the manipulations for  $x$ , as for  $y$  they are essentially the same.

Denote the list representation of  $x$  as  $(X_0, Y_0), \dots, (X_p, Y_p)$  where  $p$  is the number of pieces of  $x$ . Now we iterate over all pieces of  $x$  and check whether there exists a piece with value  $Y_j = \lfloor \alpha \rfloor_\delta^*$ . If no such piece exists, we insert  $(i, \lfloor \alpha \rfloor_\delta^*)$  into the list representation at the appropriate position. Then we find the smallest integer  $j$  such that  $Y_j > \lfloor \alpha \rfloor_\delta^*$  and for all  $k \geq j$ , we increment  $X_k$  by 1. Intuitively, we are moving all pieces that are larger than  $\lfloor \alpha \rfloor_\delta^*$  one unit to the right in order to make space for the element that was just inserted.

Once we have updated  $x$  and  $y$  as described above, we simply run the static algorithm without the step in which we initialize  $x$  and  $y$ . Note that, since we assume that after each insertion  $x$  and  $y$  are still ordered and since we only insert rounded entries into  $x$  and  $y$ , we get that  $x$  and  $y$  never have more than  $O(1/\delta)$  pieces by Lemma 39. Now, since above we have set  $\delta = \epsilon/2$ , the proof of Proposition 45 implies that we obtain a solution with additive error  $\epsilon$  in time  $O(1/\epsilon^2 \log(1/\epsilon))$ . Furthermore, note that since we do not store  $x$  and  $y$  explicitly (we only store their rounded version represented by their list representations), the space usage is  $O(1/\epsilon)$ .

Finally, we note that the operation  $Delete(i)$  can be implemented similar to above by first manipulating the list representations of  $x$  and  $y$  to remove the  $i$ 'th entries from  $x$  and  $y$  and then running the static algorithm.  $\blacktriangleleft$

We remark that by storing two dynamic vectors  $x$  and  $y$  that are undergoing element insertions and deletions as described in the proof of Proposition 48, we can also efficiently maintain an approximation of their  $(\min, +)$ -convolution  $x \oplus y$  via Lemma 40.

## H Omitted Proofs

### H.1 Proof of Lemma 6

Denote the list representations of  $g$  and  $h$  as  $(x_1^g, y_1^g), \dots, (x_{p_g}^g, y_{p_g}^g)$  and  $(x_1^h, y_1^h), \dots, (x_{p_h}^h, y_{p_h}^h)$ , respectively. Recall that both list representation are stored in doubly linked lists and that the pieces of  $g$  and  $h$  are stored in a binary search tree such that for all  $x \in [0, t]$  we can evaluate  $g(x)$  and  $f(x)$  in time  $O(\log p_g)$  and  $O(\log p_h)$ , respectively.

We show how to construct each of the functions  $f_{\min}$ ,  $f_{\text{shift}}$ ,  $f_{\text{add}}$  and  $f_{\text{round}}$  by showing how to construct their list representations.

First, let us consider  $f_{\min}$ . We construct the list representation  $(x_1^{\min}, y_1^{\min}), \dots$  of  $f_{\min}$ . The intuition of our approach is that each piece of  $f_{\min}$  must start and end at one of the start or end points of the pieces of  $g$  and  $h$ . Thus, we will evaluate the function  $\min\{g, h\}$  at all points  $x_i^g$  and  $x_i^h$  and set  $f_{\min}$  accordingly; then if  $f_{\min}$  contains multiple pieces with the same  $y_i^{\min}$ -value, we will remove these duplicate pieces. More concretely, we consider the set  $X = \{x_1^g, \dots, x_{p_g}^g, x_1^h, \dots, x_{p_h}^h\}$  and order it from small to large. Now we set  $x_i^{\min}$  to the  $i$ 'th smallest element in  $X$  for all  $i = 1, \dots, p_g + p_h$ . Observe that on the interval  $[x_{i-1}^{\min}, x_i^{\min})$ ,  $f_{\min}$  must take the value  $\min\{g(x_{i-1}^{\min}), h(x_{i-1}^{\min})\}$ . Therefore, we set  $y_i^{\min} = \min\{g(x_{i-1}^{\min}), h(x_{i-1}^{\min})\}$ . This gives an initial list representation of  $f_{\min}$ . Then we ‘‘prune’’ the list representation of  $f_{\min}$ , i.e., we iterate over all pairs  $(x_i^{\min}, y_i^{\min})$  in increasing order of  $i$  and if  $y_{i-1}^{\min} = y_i^{\min}$  then we remove the pair  $(x_{i-1}^{\min}, y_{i-1}^{\min})$  from the list representation of  $f_{\min}$ . Observe that at

the end of this process, all values of  $y_i^{\min}$  are pairwise disjoint (since the functions  $g$  and  $h$  are monotone).

To see that  $f_{\min}(x) = \min\{g(x), h(x)\}$  for all  $x \in [0, t]$ , we observe that for all  $x \in X$  (where  $X$  is as in the paragraph above) we have set  $f_{\min}(x)$  correctly by construction. Furthermore, on all contiguous intervals in  $[0, t] \setminus X$ ,  $g$  and  $h$  are constant and thus  $f_{\min}$  is constant. Therefore, for all  $x \in [0, t] \setminus X$ ,  $f_{\min}(x)$  is also set correctly.

Next, we observe that  $f_{\min}$  has at most  $p_g + p_h$  pieces because  $X$  consisted of at most  $p_g + p_h$  elements and after that we only removed pieces from  $f_{\min}$ . Furthermore, ordering the elements in  $X$  can be done in time  $O((p_g + p_h) \log(p_g + p_h))$  and evaluating  $\min\{g(x_i^{\min}), h(x_i^{\min})\}$  can be done in time  $O(\log(p_g) + \log(p_h))$ . After that we only performed a single pass over the list representations of  $f_{\min}$  in time  $O(|X|) = O(p_g + p_h)$ . Therefore, it took time  $O((p_g + p_h) \log(p_g + p_h))$  to create the list representation of  $f_{\min}$ . Finally, note that to store the elements  $x_i^{\min}$  in the binary search tree, we need additional time  $O((p_g + p_h) \log(p_g + p_h))$ .

Now we observe that  $f_{\text{add}}$  can be computed similarly to  $f_{\min}$ : the function  $f_{\text{add}}$  only changes its functions values at the points in  $X$  (where  $X$  is as above). Therefore, we let  $x_i^{\text{add}}$  be the  $i$ 'th smallest element in  $X$  and set  $y_i^{\text{add}} = g(x_{i-1}^{\text{add}}) + h(x_{i-1}^{\text{add}})$ , followed by the same pruning step as above. The rest of the proof goes through as above.

Next, consider  $f_{\text{shift}}$ . We construct the list representation  $(x_1^{\text{shift}}, y_1^{\text{shift}}), \dots, (x_{p_g}^{\text{shift}}, y_{p_g}^{\text{shift}})$  of  $f_{\text{shift}}$ . For all  $i = 1, \dots, p_g$ , we set  $x_i^{\text{shift}} = x_i^g + c$  and  $y_i^{\text{shift}} = y_i^g$ . The correctness is straightforward and from the construction it is evident that there are only  $p_g$  pieces and that everything can be done in time  $O(p_g \log(p_g))$  (since we still need to construct the binary tree for the pieces of  $f_{\text{shift}}$ ).

Finally, us consider  $f_{\text{round}}$ . As before, we construct the list representation of  $f_{\text{round}}$ ,  $(x_1^{\text{round}}, y_1^{\text{round}}), \dots, (x_{p_g}^{\text{round}}, y_{p_g}^{\text{round}})$ . For all  $i = 1, \dots, p_g$ , we set  $x_i^{\text{round}} = x_i^g$  and  $y_i^{\text{round}} = \lceil y_i^g \rceil_{1+\delta}$ . After that, we perform the same pruning step as in the construction of  $f_{\min}$ . Since  $g$  takes values in  $W_\infty = \{0\} \cup [1, W] \cup \{+\infty\}$  and  $g$  is monotone,  $f_{\text{round}}$  can take at most  $2 + \lceil \log_{1+\delta}(W) \rceil$  different values. Again, the running time bound stems from the fact that we have to construct the binary search tree for the pieces of  $f_{\text{round}}$ .

## H.2 Proof of Lemma 7

Let  $(x_1^s, y_1^s), \dots, (x_{p_s}^s, y_{p_s}^s)$  be the list representation of  $f_s$  for  $s = 1, 2$ , where  $p_s \leq p$  is the number of pieces of  $f_s$ . We create pairs  $(y_i^1, y_j^2)$  for all  $(i, j) \in \{1, \dots, p_1\} \times \{1, \dots, p_2\}$ , and order them such that  $y_i^1 + y_j^2$  becomes monotonically increasing. We iterate over all pairs in this order, and in each iteration we set the function value  $f(x)$  for some  $x$ -values to  $y := y_i^1 + y_j^2$ , where  $(y_i^1, y_j^2)$  is the pair considered during the iteration. Here, we start with large  $x$ -values (at which  $f$  takes the smallest values) and keep on decreasing  $x$  (and the function values increase); in other words, we construct  $f$  on its domain  $[0, t]$  from right to left. More concretely, let  $x_{\max}$  denote the highest  $x$ -value for which we did not yet set a function value (or  $-\infty$  if all function values have been set). Let  $x' = x_{i-1}^1 + x_{j-1}^2$ . We set the function values for all  $x \in [x', x_{\max}]$  to  $y$  and then set  $x_{\max} = x'$ . For each such new piece of  $f$ , we store that we combined the indices  $i$  and  $j$  of the pieces that we used from  $f_1$  and from  $f_2$ . Then we proceed with next iteration until all function values have been set.

The following two statements show that this procedure is correct.

1. Each  $x$  is assigned a function value that is at most the correct value  $f(x)$ . To see this let  $x \in [0, t]$  and recall that

$$f(x) = \min_{\bar{x} \in [0, x]} f_1(\bar{x}) + f_2(x - \bar{x}) .$$

Let  $\bar{x}^*$  denote the value of  $\bar{x}$  that attains the minimum in the above expression, and let  $i^*$  and  $j^*$  denote the indices of the pieces that  $\bar{x}^*$  and  $x - \bar{x}^*$  fall into, w.r.t. the list representations of  $f_1$  and  $f_2$ , respectively. This means  $\bar{x}^* \in [x_{i^*-1}^1, x_{i^*}^1)$  and  $x - \bar{x}^* \in [x_{j^*-1}^2, x_{j^*}^2)$ . Hence,  $x \geq x' = x_{i^*-1}^1 + x_{j^*-1}^2$ . Therefore, either in the iteration for the pair  $(y_{i^*}^1, y_{j^*}^2)$  or before, the procedure assigns a function value to  $x$ . Because the procedure assigns function-values in monotonically increasing fashion we are guaranteed that the function value that is assigned is at most the correct value.

2. **The function value  $y$  that is assigned is at least the correct value  $f(x)$ .** Suppose that during some iteration we assign the function value  $y = y_i^1 + y_j^2$  to  $x \in [x', \dots, x_{\max})$ , where  $x' = x_{i-1}^1 + x_{j-1}^2$ . We have

$$\begin{aligned}
 f(x) &= \min_{\bar{x} \in [0, x]} f_1(\bar{x}) + f_2(x - \bar{x}) && \text{(definition)} \\
 &\leq f_1(x_{i-1}^1) + f_2(x - x_{i-1}^1) && \text{(consider } \bar{x} = x_{i-1}^1) \\
 &\leq f_1(x_{i-1}^1) + f_2(x' - x_{i-1}^1) && (x' \leq x, f_2 \text{ monotonically decreasing)} \\
 &= f_1(x_{i-1}^1) + f_2(x_{j-1}^2) \\
 &= y_i^1 + y_j^2 && (y_i^1 = f_1(x_{i-1}^1) \text{ and } y_j^2 = f_2(x_{j-1}^2)) \\
 &= y .
 \end{aligned}$$

Hence, the assigned value is at least  $f(x)$ .

Observe that we can implement the above procedure in time  $O(p^2 \log p)$ : We first sort the at most  $p^2$  pairs in time  $O(p^2 \log p)$ . Then every iteration can be executed in constant time because setting the function values for  $x \in [x', x_{\max})$  to  $y$  can be performed by adding the pair  $(x_{\max}, y)$  to the list-representation of  $f$  and updating  $x_{\max}$  to  $x'$  takes time  $O(1)$ .

Finally, suppose we already computed  $f$  and, given  $x \in [0, t]$ , we shall return a value  $\bar{x}^* \in [0, t]$  such that  $f(x) = f_1(\bar{x}^*) + f_2(x - \bar{x}^*)$ . First, let  $(x_1, y_1), \dots, (x_p, y_p)$  denote the list representation of  $f$ . Then we can determine the piece  $\ell$  of  $f$  such that  $x \in [x_\ell, x_{\ell+1})$  in time  $O(\log p)$  since we store the pairs  $(x_i, y_i)$  of  $f$  in a binary search tree. Recall that for each piece of  $f$ , we stored the indices  $i$  and  $j$  of the pieces from  $f_1$  and  $f_2$  that we combined. Now observe that we have  $\bar{x}^* \in [x_{i-1}^1, x_i^1)$  and  $x - \bar{x}^* \in [x_{j-1}^2, x_j^2)$ , where  $i$  and  $j$  are such that these pieces from  $f_1$  and  $f_2$  form the corresponding piece of  $f$ . Thus, to find  $\bar{x}^*$  we can first try to set  $\bar{x}^* = x_{i-1}^1$ . If  $x - \bar{x}^* = x - x_{i-1}^1 \in [x_{j-1}^2, x_j^2)$  then we are done. Otherwise, we must have that  $x - x_{i-1}^1 \geq x_j^2$ . Thus, we have to increase the value of  $\bar{x}^*$  from  $x_{i-1}^1$  until it is large enough such that  $x - \bar{x}^* \in [x_{j-1}^2, x_j^2)$ . This can be achieved by setting  $\Delta = (x - x_{i-1}^1) - x_j^2$  and  $\bar{x}^* = x_{i-1}^1 + \Delta + \frac{1}{2} \min\{x_i^1 - (x_{i-1}^1 + \Delta), x_j^2 - x_{j-1}^2\}$ . Note that this value of  $\bar{x}^*$  can be computed in time  $O(1)$ . Thus, the total time to return  $\bar{x}^*$  is  $O(\log p)$ .

### H.3 Proof of Theorem 9

Recall that the dependency graph is a DAG. We call a vertex without any incoming edges a *leaf*. The *level* of a vertex  $u$  is the length of the longest path from a leaf to  $u$ . Note that since each node can only reach  $h$  other nodes, every vertex has level at most  $h$ .

We compute the DP bottom-up, starting at the leaves of the DAG and then recursively computing the solutions for rows  $i$  for which the solutions of  $\text{In}(i)$  have already been computed. We store the approximate solutions  $\text{ADP}(i, \cdot)$  using monotone piecewise constant functions.

We prove the theorem by induction over the level of  $i$  in the dependency graph. We show the stronger statement that for every DP row  $i$  of level  $\ell$ ,  $\text{ADP}(i, \cdot)$  is an  $\alpha^{\ell+1}$ -approximation of  $\text{DP}(i, \cdot)$ .

We start with leaf vertices (i.e., vertices of level 0). For a leaf  $i$ , we use Properties 4(b) and 4(c) to obtain that  $\tilde{\mathcal{P}}_i$  returns  $\text{ADP}(i, \cdot)$  which is a monotone piecewise constant function with at most  $p$  pieces and which is an  $\alpha$ -approximation of  $\text{DP}(i, \cdot)$ .

Next, consider a row  $i$  of level  $\ell$ . We use  $\tilde{\mathcal{P}}_i$  to compute  $\text{ADP}(i, \cdot) = \tilde{\mathcal{P}}_i(\{\text{ADP}(i', \cdot) : i' \in \text{In}(i)\})$ . By induction hypothesis, all solutions  $\text{ADP}(i', \cdot)$ ,  $i' \in \text{In}(i)$ , are stored as monotone piecewise constant functions and each of them has at most  $p$  pieces. Since we apply the operations from Lemma 6 only  $O(1)$  times, the number of pieces only grows by a factor  $O(1)$ . Since we only apply the  $(\min, +)$ -convolution from Lemma 7 at most a single time, the number of pieces after the convolution is bounded by  $O(p^2)$ . Thus, we will never operate on functions with more than  $O(p^2)$  pieces. The bounds from Lemmas 6 and 7 imply that all operations to compute  $\tilde{\mathcal{P}}_i$  can be performed in time at most  $O(p^2 \log(p))$ . Furthermore, by induction hypothesis and since each  $i'$  is at level  $\ell' \leq \ell - 1$ , we know that  $\text{ADP}(i', \cdot)$  is an  $\alpha^{\ell'}$ -approximation of  $\text{DP}(i', \cdot)$ . Using Properties (3) and 4(a), we get that  $\text{ADP}(i, \cdot)$  is an  $\alpha^{\ell+1}$ -approximation of  $\text{DP}(i, \cdot)$ .

The theorem's approximation guarantee follows from Property (2) which implies that  $\ell \leq h$  for all DP rows  $i$  in the dependency graph. Furthermore, above we argued that each solution  $\text{ADP}(i, \cdot)$  can be computed in time  $O(p^2 \log(p))$  which gives a total running time of  $O(|I| \cdot p^2 \log(p))$ .

#### H.4 Proof of Theorem 10

Consider a row  $i$  for which  $\text{DP}(i, \cdot)$  changes. Note that we only have to compute DP solutions for rows  $i'$  which are reachable from  $i$  in the dependency graph. Since we assume that the dependency graph is a DAG and  $\text{Reach}(i) \leq h$  for all rows  $i$ , there can be at most  $h$  such rows. In the proof of Theorem 9 we argued that each solution  $\text{ADP}(i, \cdot)$  can be computed in time  $O(p^2 \log(p))$ . This gives the proof of the theorem.

#### H.5 Property of the Racke Tree

Let  $G = (V_G, E_G)$  be an undirected graph and let  $T = (V_T, E_T)$  be a Racke tree for  $G$ . We prove that  $\text{mincut}_T(A, B) \geq \text{mincut}_G(A, B)$  by showing that for any set of vertices  $S_T \subseteq V_T$ , it holds that  $\text{cap}_T(S_T) \geq \text{cap}_G(S)$  where  $S \subseteq V_G$  is the set of leaf vertices in  $S_T$ .

Let  $S_T \subseteq V_T$  and consider the cut  $(S_T, \bar{S}_T)$  in  $T$ . We use  $S$  to denote the restriction of  $S_T$  to the leaf vertices and observe that  $(S, \bar{S})$  forms a cut in  $G$  as well. Then:

$$\begin{aligned}
\text{cap}_T(S_T) &= \sum_{(x_t, y_t) \in S_T \times \bar{S}_T} \text{cap}_T(x_t, y_t) && \text{(definition of } \text{cap}_T(S_T)) \\
&= \sum_{(x_t, y_t) \in S_T \times \bar{S}_T} \text{cap}_G(V_{x_t} \cap V_{y_t}) && \text{(definition of tree edge capacity)} \\
&= \sum_{(x_t, y_t) \in S_T \times \bar{S}_T} \sum_{(x, y) \in V_{x_t} \times \bar{V}_{y_t}} \text{cap}_G(x, y) && \text{(w.l.o.g. assume } V_{x_t} \subseteq V_{y_t}) \\
&= \sum_{\{x, y\} \in E_G} \text{cap}_G(x, y) \sum_{(x_t, y_t) \in S_T \times \bar{S}_T} \mathbb{1}\{x \in V_{x_t} \wedge y \in \bar{V}_{y_t}\} && \text{(change order of summation)} \\
&\geq \sum_{(x, y) \in S \times \bar{S}} \text{cap}_G(x, y) = \text{cap}_G(S) .
\end{aligned}$$

Here the inequality follows because a pair  $(x, y) \in S \times \bar{S}$  whose capacity is counted on the right hand side corresponds to a graph edge  $\{x, y\} \in E_G$  (between  $x \in S$  and  $y \in \bar{S}$ ). This



graph edge contributes to the capacity on every edge of the  $x$ - $y$  path in  $T$ . One of these edges must be cut by  $S_T$ , i.e.,  $\mathbb{1}\{x \in V_{x_t} \wedge y \in \bar{V}_{x_t}\} = 1$  for this tree edge. Hence, its capacity is also counted on the left hand side.

## H.6 Proof of Lemma 18

The lower bound is immediate since  $\tilde{\mathcal{P}}_i$  is an  $\alpha$ -approximation of  $\mathcal{P}_i$ . For the upper bound we use induction over  $\ell$ . For  $\ell = 0$  observe that

$$\begin{aligned} \text{ADP}(i, \cdot) &= \tilde{\mathcal{P}}_i(\{\text{ADP}(i') : i' \in \text{In}(i)\}) && \text{(definition)} \\ &\leq \alpha \mathcal{P}_i(\{\text{ADP}(i') : i' \in \text{In}(i)\}) && (\tilde{\mathcal{P}}_i \text{ is } \alpha\text{-approximate}) \\ &= \alpha \mathcal{P}_i(\emptyset) && (v_i \text{ is a leaf}) \\ &= \alpha \text{DP}(i) && \text{(the DP is okay-behaved)} \end{aligned}$$

For  $\ell > 0$  we have

$$\begin{aligned} \text{ADP}(i, \cdot) &= \tilde{\mathcal{P}}_i(\{\text{ADP}(i') : i' \in \text{In}(i)\}) && \text{(definition)} \\ &\leq \alpha \mathcal{P}_i(\{\text{ADP}(i') : i' \in \text{In}(i)\}) && (\tilde{\mathcal{P}}_i \text{ is } \alpha\text{-approximate}) \\ &= \alpha \mathcal{P}_i(\{\alpha^\ell \text{DP}(i', \cdot) : i' \in \text{In}(i)\}) && \text{(induction hypothesis)} \\ &= \alpha^{\ell+1} \text{DP}(i, \cdot) && \text{(the DP is okay-behaved)} \end{aligned}$$

Here the induction hypothesis exploits the fact that all  $i' \in \text{In}(i)$ , have level strictly less than  $\ell$  in the dependency graph.

## H.7 Proof of Lemma 19

The claim about the approximation guarantee follows immediately from Lemma 18 and the fact that the root has level at most  $h$  (since the longest leaf-root path in the dependency tree has length  $h$ ). To obtain running time  $O(|V_T| \cdot t)$ , we compute the solutions  $\text{ADP}(v_1), \dots, \text{ADP}(v_n)$  in this order, i.e., based on the topological ordering of the dependency DAG. Then by assumption on the ordering of the rows  $i$  and since all  $\tilde{\mathcal{P}}_i$  can be computed in time  $t$ , the lemma follows.

## H.8 Proof of Lemma 20

Suppose the inserted or deleted edge is incident upon a vertex  $i$ . Since the DPs we consider are well-behaved, we only need to recompute DP solutions for those vertices  $j$  such that there exists a directed path from  $i$  to  $j$ ,  $j \geq i$ , in the dependency graph. By construction of the dependency graph, there can be at most  $h$  such vertices (since the longest leaf-root path in the dependency graph has length  $h$ ). Therefore, we can recompute all of these solutions in time  $O(h \cdot t)$ . After we finished the recomputation, the guarantees on the approximation ratio are implied by Lemma 19.

## H.9 Proof of Lemma 21

We only prove the case if all functions  $f_i$  are monotonically decreasing. The case for monotonically increasing functions is analogous. Let  $\mathcal{P}$  denote the set of all pieces in functions  $f_i$ . Consider a piece  $p \in \mathcal{P}$  that starts at  $t_1$  ends at  $t_2$  and has value  $\alpha$ . We construct a piece-wise constant function  $f_p : [0, t] \rightarrow W_\infty$  with two pieces that has value  $\infty$  on  $[0, t_1)$ , and value  $\alpha$  on  $[t_1, t]$  (this means we extend the piece from  $t_2$  to  $t$ ).

Because the functions are monotonically decreasing we can rewrite  $f_{\min}$  as a minimum of the piece-functions  $f_p$ , i.e.,

$$f_{\min}(x) = \min_{p \in \mathcal{P}} f_p(x) .$$

We now sort all pieces in  $\mathcal{P}$  by their start-point. By processing the pieces in sorting order we can build the result function step-by-step. Let  $f_{r-1}$  denote the piece-wise constant function encoding the minimum over the first  $r-1$  pieces. In order to compute  $f_r$  we have to compare the last piece of  $f_{r-1}$  to the  $r$ -th piece  $p_r$  in  $\mathcal{P}$ . If the value of  $p_r$  is higher than  $f_{r-1}(\infty)$  (the value of the last piece in  $f_{r-1}$ ) we ignore the piece  $p_r$ . Otherwise, we end the current last piece of  $f_{r-1}$  at the start time  $t_r$  of piece  $p_r$  and add the piece  $p_r$  with its start time, its value, and an end time of  $t$ . The running time is dominated by sorting the pieces and inserting them into a binary search tree when adding them to the result function.

### H.10 Proof of Lemma 22

We can assume w.l.o.g. that  $f_2$  is monotonically decreasing (this follows from the symmetry of  $(\min, +)$ -convolution). Now the lemma is implied by the following computation, where in the third step we use the monotonicity of  $f_2$ , i.e., we use that  $f_2(x') \leq f_2(x)$  for all  $x' \geq x$ :

$$\begin{aligned} f(x') &= \min_{\bar{x} \in [0, x']} f_1(\bar{x}) + f_2(x' - \bar{x}) \\ &\leq \min_{\bar{x} \in [0, x]} f_1(\bar{x}) + f_2(x' - \bar{x}) \\ &\leq \min_{\bar{x} \in [0, x]} f_1(\bar{x}) + f_2(x - \bar{x}) \\ &= f(x). \end{aligned}$$