



HAL
open science

Approximation Strategies for Generalized Binary Search in Weighted Trees

Dariusz Dereniowski, Adrian Kosowski, Przemyslaw Uznanski, Mengchuan
Zou

► **To cite this version:**

Dariusz Dereniowski, Adrian Kosowski, Przemyslaw Uznanski, Mengchuan Zou. Approximation Strategies for Generalized Binary Search in Weighted Trees . ICALP 2017 - 44th International Colloquium on Automata, Languages, and Programming, 2017, Warsaw, Poland. pp.84:1–84:14, 10.4230/LIPIcs.ICALP.2017.84 . hal-01470935

HAL Id: hal-01470935

<https://inria.hal.science/hal-01470935v1>

Submitted on 26 Feb 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Approximation Strategies for Generalized Binary Search in Weighted Trees

Dariusz Dereniowski¹, Adrian Kosowski², Przemysław Uznański³, and Mengchuan Zou²

¹Faculty of Electronics, Telecommunications and Informatics,
Gdańsk University of Technology, Poland

²Inria Paris and IRIF, Université Paris Diderot, France

³Department of Computer Science, ETH Zürich, Switzerland

Abstract

We consider the following generalization of the binary search problem. A search strategy is required to locate an unknown target node t in a given tree T . Upon querying a node v of the tree, the strategy receives as a reply an indication of the connected component of $T \setminus \{v\}$ containing the target t . The cost of querying each node is given by a known non-negative weight function, and the considered objective is to minimize the total query cost for a worst-case choice of the target.

Designing an optimal strategy for a weighted tree search instance is known to be strongly NP-hard, in contrast to the unweighted variant of the problem which can be solved optimally in linear time. Here, we show that weighted tree search admits a quasi-polynomial time approximation scheme (QPTAS): for any $0 < \varepsilon < 1$, there exists a $(1+\varepsilon)$ -approximation strategy with a computation time of $n^{O(\log n/\varepsilon^2)}$. Thus, the problem is not APX-hard, unless $NP \subseteq DTIME(n^{O(\log n)})$. By applying a generic reduction, we obtain as a corollary that the studied problem admits a polynomial-time $O(\sqrt{\log n})$ -approximation. This improves previous $\tilde{O}(\log n)$ -approximation approaches, where the \tilde{O} -notation disregards $O(\text{poly log log } n)$ -factors.

Key Words: Approximation Algorithm; Adaptive Algorithm; Graph Search; Binary Search; Vertex Ranking; Trees

1 Introduction

In this work we consider a generalization of the fundamental problem of searching for an element in a sorted array. This problem can be seen, using graph-theoretic terms, as a problem of searching for a target node in a path, where each query reveals on which ‘side’ of the queried node the target node lies. The generalization we study is two-fold: a more general structure of a tree is considered and we assume non-uniform query times. Thus, our problem can be stated as follows. Given a node-weighted input tree T (in which the query time of a node is provided as its weight), design a search strategy (sometimes called a decision tree) that locates a hidden *target node* x by asking *queries*. Each query selects a node v in T and after the time that equals the weight of the selected node, a reply is given: the reply is either ‘yes’ which implies that v is the target node and thus the search terminates, or it is ‘no’ in which case the search strategy receives the edge outgoing from v that belongs to the shortest path between u and v . The goal

is to design a search strategy that locates the target node and minimizes the search time in the worst case.

The vertex search problem is more general than its ‘edge variant’ that has been more extensively studied. In the latter problem one selects an edge e of an edge-weighted tree $T = (V, E, w)$ in a query and learns in which of the two components of $T - e$ the target node is located. Indeed, this edge variant can be reduced to our problem as follows: first assign a ‘large’ weight to each node of T (for example, one plus the sum of the weights of all edges in the graph) and then subdivide each edge e of T giving to the new node the weight of the original edge, $w(e)$. It is apparent that an optimal search strategy for the new node-weighted tree should never query the nodes with large weights, thus immediately providing a search strategy for the edge variant of T .

We also point out that the considered problem, as well as the edge variant, being quite fundamental, were historically introduced several times under different names: minimum height elimination trees [31], ordered colourings [19], node and edge rankings [16], tree-depth [29] or LIFO-search [14].

Table 1.1 summarizes the complexity status of the node-query model (in case of unweighted paths in both cases the solution is the classical binary search algorithm) and places our result in the general context.

Table 1.1: Computational complexity of the search problem in different graph classes, including our results for weighted trees. Completeness results refer to the decision version of the problem.

<i>Graph class</i>	<i>Unweighted</i>	<i>Weighted</i>
Paths:	exact in $O(n)$ time	exact in $O(n^2)$ time [6] strongly NP-complete [11]
Trees:	exact in $O(n)$ time [30, 33]	$(1 + \varepsilon)$ -approx. in $n^{O(\log n/\varepsilon)}$ time (Thm. 3.3) $O(\sqrt{\log n})$ -approx. in poly-time (Thm. 3.4)
Undirected:	exact in $n^{O(\log n)}$ time [12] $O(\log n)$ -approx. in poly-time [12]	PSPACE-complete [12] $O(\log n)$ -approx. in poly-time [12]
Directed:	PSPACE-complete [12]	PSPACE-complete [12]

1.1 State-of-the-Art

In this work we focus on the worst case search time for a given input graph and we only remark that other optimization criteria has been also considered [5, 20, 21, 35]. For other closely related models and corresponding results see e.g. [1, 15, 23, 25, 34].

The node-query model. An optimal search strategy can be computed in linear-time for an unweighted tree [30, 33]. The number of queries performed in the worst case may vary from being constant (for a star one query is enough) to being at most $\log_2 n$ for any tree [30] (by always querying a node that halves the search space). Several following results have been obtained in [12]. First, it turns out that $\log_2 n$ queries are always sufficient for general simple graphs and this implies a $O(m^{\log_2 n} n^2 \log n)$ -time optimal algorithm for arbitrary unweighted graphs. The algorithm which performs $\log_2 n$ queries also serves as a $O(\log n)$ -approximation algorithm, also for the weighted version of the problem. (We remark that in the weighted case, the algorithms

in [12] sometimes have an approximation ratio of $\Theta(\log n)$, even in the tree scenario we study in this work.) On the other hand, it is shown in the same work that an optimal algorithm (for unweighted case) with a running time of $O(n^{o(\log n)})$ would be in contradiction with the Exponential-Time-Hypothesis, and for $\varepsilon > 0$, $O(m^{(1-\varepsilon)\log n})$ would be in contradiction with the Strong Exponential-Time-Hypothesis. When weighted graphs are considered, the problem becomes PSPACE-complete. Also, a generalization to directed graphs also turns out to be PSPACE-complete.

We also refer the interested reader to further works that consider a probabilistic version of the problem, where the answer to a query is correct with some probability $p > \frac{1}{2}$ [3, 12, 13, 18]. In particular, for any $p > \frac{1}{2}$ and any undirected unweighted graph, a search strategy can be computed that finds the target node with probability $1 - \delta$ using $(1 - \delta) \frac{\log_2 n}{1 - H(p)} + o(\log n) + O(\log^2 \frac{1}{\delta})$ queries in expectation, where $H(p) = -p \log_2 p - (1 - p) \log_2 (1 - p)$ is the entropy function. See [32] for a model in which a fixed number of queries can be answered incorrectly during a binary search.

The edge-query model. In the case of unweighted trees, an optimal search strategy can be computed in linear time [24, 28]. (See [9] for a correspondence between edge rankings and the searching problem.) The problem of computational complexity for weighted trees attracted a lot of attention. On the negative side, it has been proved that it is strongly NP-hard to compute an optimal search strategy [8] for bounded diameter trees, which has been improved by showing hardness for several specific topologies: trees of diameter at most 6, trees of degree at most 3 [6] and spiders [7] (trees having at most one node of degree greater than two). On the other hand, polynomial-time algorithms exist for weighted trees of diameter at most 5 and weighted paths [6]. We note that for weighted paths there exists a linear-time but approximate solution given in [20]. For approximate polynomial-time solutions, a simple $O(\log n)$ -approximation has been given in [8] and a $O(\log n / \log \log \log n)$ -approximate solution is given in [6]. Then, the best known approximation ratio has been further improved to $O(\log n / \log \log n)$ in [7].

Some bounds on the number of queries for unweighted trees have been developed. Observe that an optimal search strategy needs to perform at least $\log_2 n$ queries in the worst case. However, there exist trees of maximum degree Δ that require $\Delta \log_{\Delta+1} n$ queries [2]. On the other hand, $\Theta(\Delta \log n)$ queries are always sufficient for each tree [2], which has been improved to $(\Delta + 1) \log_{\Delta} n$ [22], $\Delta \log_{\Delta} n$ [10] and $1 + \frac{\Delta-1}{\log_2(\Delta+1)-1} \log_2 n$ [12].

Searching partial orders. The problem of searching a partial order with uniform query times is NP-complete even for partial orders with maximum element and bounded height Hasse diagram [4, 9]. For some algorithmic solutions for random partial orders see [4]. For a given partial order P with maximum element, an optimal solution can be obtained by computing a branching B (a directed spanning tree with one target) of the directed graph representing P and then finding a search strategy for the branching, as any search strategy for B also provides a feasible search for P [9]. Since computing an optimal search strategy for B can be done efficiently (through the equivalence to the edge-query model), finding the right branching is a challenge. This approach has been used in [9] to obtain an $O(\log n / \log \log n)$ -approximation polynomial time algorithm for partial orders with a maximum element.

We remark that searching a partial order with a maximum element or with a minimum element are essentially quite different. For the latter case a linear-time algorithm with additive error of 1 has been given in [30]. As observed in [9], the problem of searching in tree-like partial orders with a maximum element (which corresponds to the edge-query model in trees) is equivalent to the edge ranking problem.

1.2 Organization of the Paper

The aim of Section 2 is to give the necessary notation and a formal statement of the problem (Sections 2.1 and 2.2) and to provide two different but equivalent problem formulations that will be more convenient for our analysis. As opposed to the classical problem formulation in which a strategy is seen as a *decision tree*, Section 2.3 restates the problem in such a way that with each vertex v of the input tree we associate a sequence of vertices that need to be iteratively queried when v is the root of the current subtree that contains the target node. In Section 2.4 we extend this approach by associating with each vertex a sequence of not only vertices to be queried but also time points of the queries.

The latter problem formulation is suitable for a dynamic programming algorithm provided in Section 3.1. In this section we introduce an auxiliary, slightly modified measure of the cost of a search strategy. First we provide a quasi-polynomial time dynamic programming scheme that provides an arbitrarily good approximation of the output search strategy with respect to this modified cost (the analysis is deferred to Section 4), and then we prove that the new measure is sufficiently close to the original one (the analysis is deferred to Section 5). These two facts provide the quasi-polynomial time scheme for the tree search problem, achieving a $(1 + \varepsilon)$ -approximation with a computation time of $n^{O(\log n/\varepsilon^2)}$, for any $0 < \varepsilon < 1$.

In Section 3.2 we observe how to use the above algorithm to derive a polynomial-time $O(\sqrt{\log n})$ -approximation algorithm for the tree search problem. This is done by a divide and conquer approach: a sufficiently small subtree T^* of the input tree T is first computed so that the quasi-polynomial time algorithm runs in polynomial (in the size of T) time for T^* . This decomposes the problem: having a search strategy for T^* , the search strategies for $T - T^*$ are computed recursively. Details of the approach are provided in Section 6.

2 Preliminaries

2.1 Notation and Query Model

We now recall the problem of searching of an unknown target node x by performing queries on the vertices of a given node-weighted rooted tree $T = (V, E, w)$ with weight function $w: V \rightarrow \mathbb{R}_+$. Each *query* selects one vertex v of T and after $w(v)$ time units receives an answer: either the query returns *true*, meaning that $x = v$, or it returns a neighbor u of v which lies closer to the target x than v . Since we assume that the queried graph T is a tree, such a neighbor u is unique and is equivalently described as the unique neighbor of v belonging to the same connected component of $T \setminus \{v\}$ as x .

All trees we consider are rooted. Given a tree T , the root is denoted by $r(T)$. For a node $v \in V$, we denote by T_v the subtree of T rooted at v . For any subset $V' \subseteq V$ (respectively, $E' \subseteq E$) we denote by $T[V']$ (resp., $T[E']$) the minimal subtree of T containing all nodes from V' (resp., all edges from E'). For $v \in V$, $N(v)$ is the set of neighbors of v in T .

For $U \subseteq V$ and a target node $x \notin U$, there exists a unique maximal subtree of $T \setminus U$ that contains x ; we will denote this subtree by $T \langle U, x \rangle$.

We denote $|V| = n$. We will assume w.l.o.g. that the maximum weight of a vertex is normalized to 1. (This normalization is immediately obtained by a proportional scaling of all units of cost.) We will also assume w.l.o.g. that the weight function satisfies the following *star condition*:

$$\text{for all } v \in V, w(v) \leq \sum_{u \in N(v)} w(u).$$

Observe that if this condition is not fulfilled, i.e., for some vertex v will have $w(v) > \sum_{u \in N(v)} w(u)$,

then vertex v will never be queried by any optimal strategy in v , since a query to v can then be replaced by a sequence of queries to all neighbors of v , obtaining not less information at strictly smaller cost. In general, given an instance which does not satisfy the star condition, we enforce it by performing all necessary weight replacements $w(v) \leftarrow \min\{w(v), \sum_{u \in N(v)} w(u)\}$, for $v \in V$.

For $a, \omega \in \mathbb{R}_{\geq 0}$, we denote the rounding of a down (up) to the nearest multiple of ω as $\lfloor a \rfloor_\omega = \omega \lfloor a/\omega \rfloor$ and $\lceil a \rceil_\omega = \omega \lceil a/\omega \rceil$, respectively.

2.2 Definition of a Search Strategy

A *search strategy* \mathcal{A} for a rooted tree $T = (V, E, w)$ is an adaptive algorithm which defines successive queries to the tree, based on responses to previous queries, with the objective of locating the target vertex in a finite number of steps. Note that search strategies can be seen as decision trees in which each node represents a subset of vertices of T that contains x , with leaves representing singletons consisting of x .

Let $\mathbf{Q}_{\mathcal{A}}(T, x)$ be the time-ordering (sequence) of queries performed by strategy \mathcal{A} on tree T to find a target vertex x , with $\mathbf{Q}_{\mathcal{A},i}(T, x)$ denoting the i -th queried vertex in this time ordering, $1 \leq i \leq |\mathbf{Q}_{\mathcal{A}}(T, x)|$.

We denote by

$$\text{COST}_{\mathcal{A}}(T, x) = \sum_{i=1}^{|\mathbf{Q}_{\mathcal{A}}(T, x)|} w(\mathbf{Q}_{\mathcal{A},i}(T, x))$$

the sum of weights of all vertices queried by \mathcal{A} with x being the target node, i.e., the time after which \mathcal{A} finishes. Let

$$\text{COST}_{\mathcal{A}}(T) = \max_{x \in V} \text{COST}_{\mathcal{A}}(T, x)$$

be the *cost* of \mathcal{A} . We define the *cost* of T to be

$$\text{OPT}(T) = \min\{\text{COST}_{\mathcal{A}}(T) \mid \mathcal{A} \text{ is a search strategy for } T\}.$$

We say that a search strategy is *optimal* for T if its cost equals $\text{OPT}(T)$.

As a consequence of normalization and the star condition, we have the following bound.

Observation 2.1. *For any tree T , we have $1 \leq \text{OPT}(T) \leq \lceil \log_2 n \rceil$.*

Proof. By the star condition, considering any vertex $v \in V$ as the target, we trivially have

$$\text{OPT}(T) \geq \inf_{\mathcal{A}} \text{COST}_{\mathcal{A}}(T, v) \geq \inf_{\mathcal{A}} \text{COST}_{\mathcal{A}}(T[\{v\} \cup N(v)], v) \geq w(v).$$

Thus, $\text{OPT}(T) \geq \max_{v \in V} w(v) = 1$, which gives the first inequality.

For the second inequality, we observe that applying to tree T the optimal search strategy for unweighted trees, we can locate the target in at most $\lceil \log_2 n \rceil$ queries (cf. e.g. [19, 30]). Since the cost of each query is at most 1, the claim follows. \square

We also introduce the following notation. If the first $|U|$ queried vertices by a search strategy \mathcal{A} are exactly the vertices in U , $U = \{\mathbf{Q}_{\mathcal{A},i}(T, x) : 1 \leq i \leq |U|\}$, then we say that \mathcal{A} *reaches* $T\langle U, x \rangle$ *through* U , and $w(U)$ is the *cost of reaching* $T\langle U, x \rangle$ *by* \mathcal{A} . We also say that we receive an ‘up’ reply to a query to a vertex v if the root of the tree remaining to be searched remains unchanged by the query, i.e., $r(T\langle U, x \rangle) = r(T\langle U \cup \{v\}, x \rangle)$, and we call the reply a ‘down’ reply when the root of the remaining tree changes, i.e., $r(T\langle U, x \rangle) \neq r(T\langle U \cup \{v\}, x \rangle)$. Without loss of generality, after having performed a sequence of queries U , we can assume that the tree $T\langle U, x \rangle$ is known to the strategy.

2.3 Query Sequences and Stable Strategies

By a slight abuse of notation, we will call a search strategy *polynomial-time* if it can be implemented using a dynamic (adaptive) algorithm which computes the next queried vertex in polynomial time.

We give most of our attention herein to search strategies in trees which admit a natural (non-adaptive, polynomial-space) representation called a *query sequence assignment*. Formally, for a rooted tree T , the *query sequence assignment* S is a function $S : V \rightarrow V^*$, which assigns to each vertex $v \in V$ an ordered sequence of vertices $S(v)$, known as the *query sequence* of v . The query sequence assignment directly induces a strategy \mathcal{A}_S , presented as Algorithm 2.1. Intuitively, the strategy processes successive queries from the sequence $S(v)$, where v is the root vertex of the current search tree, $v = r(T\langle U, x \rangle)$, where U is the set of queries performed so far. This processing is performed in such a way that the strategy iteratively takes the first vertex in $S(v)$ that belongs to $T\langle U, x \rangle$ and queries it. As soon as the root of the search tree changes, the procedure starts processing queries from the sequence of the new root, which belong to the remaining search tree. The procedure terminates as soon as $T\langle U, x \rangle$ has been reduced to a single vertex, which is necessarily the target x .

Algorithm 2.1 Search strategy \mathcal{A}_S for a query sequence assignment S

```

1:  $v \leftarrow r(T)$  // stores current root
2:  $U \leftarrow \emptyset$ 
3: while  $|T\langle U, x \rangle| > 1$  do
4:   for  $u \in S(v)$  do
5:     if  $u \in T\langle U, x \rangle$  then //  $u$  is the first vertex in  $S(v)$  that belongs to  $T\langle U, x \rangle$ 
6:       QUERYVERTEX( $u$ )
7:        $U \leftarrow U \cup \{u\}$ 
8:     if  $v \neq r(T\langle U, x \rangle)$  then // query reply is ‘down’
9:        $v \leftarrow r(T\langle U, x \rangle)$ 
10:    break // for loop

```

In what follows, in order to show that our approximation strategies are polynomial-time, we will confine ourselves to presenting a polynomial-time algorithm which outputs an appropriate sequence assignment.

A sequence assignment is called *stable* if the replacement of line 9 in Algorithm 2.1 by any assignment of the form $v \leftarrow v''$, where v'' is an arbitrary vertex which is promised to lie on the path from $r(T\langle U, x \rangle)$ to the target x , always results in a strategy which performs a (not necessarily strict) subsequence of the sequence of queries performed by the original strategy \mathcal{A}_S . Sequence assignments computed on trees with a bottom-up approach usually have the stability property; we provide a proof of stability for one of our main routines in Section 4.

Without loss of generality, we will also assume that if $v \in S(v)$, then v is the last element of $S(v)$. Indeed, when considering a subtree rooted at v , after a query to v , if v was not the target, then the root of the considered subtree will change to one of the children of v , hence any subsequent elements of $S(v)$ may be removed without changing the strategy.

2.4 Strategies Based on Consistent Schedules

Intuitively, we may represent search strategies by a schedule consisting of some number of jobs, with each job being associated to querying a node in the tree (cf. e.g. [17, 26, 27]). Each job has a fixed processing time, which is set to the weight of a node. Formally, in this work we will refer to the schedule \hat{S} only in the very precise context of search strategies \mathcal{A}_S based on some query sequence assignment S . The *schedule assignment* \hat{S} is the following extension of the sequence

assignment S , which additionally encodes the starting time of search query job. If the query sequence S of a node v is of the form $S(v) = (v_1, \dots, v_k)$, $k = |S(v)|$, then the corresponding schedule for v will be given as $\hat{S}(v) = ((v_1, t_1), \dots, (v_k, t_k))$, with $t_i \in \mathbb{R}_{\geq 0}$ denoting the starting time of the query for v_i . We will call $\hat{S}(v)$ the *schedule of node v* . We will call a schedule assignment \hat{S} *consistent* with respect to search in a given tree T if the following conditions are fulfilled:

- (i) No two jobs in the schedule of a node overlap: for all $v \in V$, for two distinct jobs $(u_1, t_1), (u_2, t_2) \in \hat{S}(v)$, we have $[[t_1, t_1 + w(u_1)] \cap [t_2, t_2 + w(u_2)]] = 0$.
- (ii) If v is the parent of v' in T and $(u, t) \in \hat{S}(v')$, then we either also have $(u, t) \in \hat{S}(v)$, or the job $(v, t_v) \in \hat{S}(v)$ completes before the start of job (u, t) : $t_v + w(v) \leq t$.

It follows directly from the definition that a consistent schedule assignment (and the underlying query sequence assignment) is uniquely determined by the collection of jobs $\{(v, t_v) : (v, t_v) \in \hat{S}(u), u \in V\}$. Note that not every vertex has to contain a query to itself in its schedule; we will occasionally write $t_v = \perp$ to denote that such a job is missing. In this case, the jobs of all children of v have to be contained in the schedule of node v .

By extension of notation for sequence assignments, we will denote a strategy following a consistent schedule assignment \hat{S} (i.e., executing the query jobs of schedule \hat{S} at the prescribed times) as $\mathcal{A}_{\hat{S}}$. We will then have:

$$\text{COST}_{\mathcal{A}_{\hat{S}}}(T) = |\hat{S}|,$$

where $|\hat{S}|$ is the *duration* of schedule assignment \hat{S} , given as:

$$|\hat{S}| = \max_{v \in V} |\hat{S}(v)|,$$

with:

$$|\hat{S}(v)| = \max_{(u, t) \in \hat{S}(v)} (t + w(u)).$$

We remark that there always exists an optimal search strategy which is based on a consistent schedule. By a well-known characterization (cf. e.g. [8]), tree T satisfies $\text{OPT}(T) = \tau \in \mathbb{R}$ if and only if there exists an assignment $I : V \rightarrow \mathcal{I}_{\tau}$ of intervals of time to nodes before deadline τ , $\mathcal{I}_{\tau} = \{[a, b] : 0 \leq a < b \leq \tau\}$, such that $|I(v)| = w(v)$ and if $|I(u) \cap I(v)| > 0$ for any pair of nodes $u, v \in V$, then the $u - v$ path in T contains a separating vertex z such that $\max I(z) \leq \min(I(u) \cup I(v))$. The corresponding schedule assignment of duration τ is obtained by adding, for each node $u \in V$, the job $(u, \min I(u))$ to the schedule of all nodes on the path from u towards the root, until a node v such that $\max I(v) \leq \min I(u)$ is encountered on this path. The consistency and correctness of the obtained schedule is immediate to verify.

Observation 2.2. *For any tree T , there exists a query sequence assignment S and a corresponding consistent schedule \hat{S} on T such that $|\hat{S}| = \text{OPT}(T)$. \square*

3 The Results

3.1 $(1 + \varepsilon)$ -Approximation in $n^{O(\log n / \varepsilon^2)}$ Time

We first present an approximation scheme for the weighted tree search problem with $n^{O(\log n)}$ running time. The main difficulty consists in obtaining a constant approximation ratio for the problem with this running time; we at once present this approximation scheme with tuned parameters, so as to achieve $(1 + \varepsilon)$ -approximation in $n^{O(\log n / \varepsilon^2)}$ time.

Our construction consists of two main building blocks. First, we design an algorithm based on a bottom-up (dynamic programming) approach, which considers exhaustively feasible sequence assignments and query schedules over a carefully restricted state space of size $n^{O(\log n)}$ for each node. The output of the algorithm provides us both with a lower bound on $\text{OPT}(T)$, and with a sequence assignment-based strategy \mathcal{A}_S for solving the tree search problem. The performance of this strategy \mathcal{A}_S is closely linked to the performance of $\text{OPT}(T)$, however, there is one type of query, namely a query on a vertex of small weight leading to a ‘down’ response, due to whose repeated occurrence the eventual cost difference between $\text{COST}_{\mathcal{A}_S}(T)$ and $\text{OPT}(T)$ may eventually become arbitrarily large. To alleviate this difficulty, we introduce an alternative measure of cost which compensates for the appearance of the disadvantageous type of query.

We start by introducing some additional notation. Let $\omega \in \mathbb{R}_+$, be an arbitrarily fixed value of weight and let $c \in \mathbb{N}$. The choice of constant $c \in \mathbb{N}$ will correspond to an approximation ratio of $(1 + \varepsilon)$ of the designed scheme for $\varepsilon = 168/c$.

We say that a query to a vertex v is a *light down query* in some strategy if $w(v) < c\omega$ and $x \in V(T_v)$, i.e., it is also a ‘down’ query, where x is the target vertex.

For any strategy \mathcal{A} , we denote by $\text{COST}_{\mathcal{A}}^{(\omega, c)}(T, x)$ its modified cost of finding target x , defined as follows. Let d_x be the number of light down queries when searching for x :

$$d_x = \left| \{i : w(Q_{\mathcal{A}, i}(T, x)) < c\omega \text{ and } x \in V(T_{Q_{\mathcal{A}, i}(T, x)})\} \right|.$$

Then, the modified cost $\text{COST}_{\mathcal{A}}^{(\omega, c)}(T, x)$ is:

$$\text{COST}_{\mathcal{A}}^{(\omega, c)}(T, x) = \text{COST}_{\mathcal{A}}(T, x) - (2c + 1)\omega d_x. \quad (3.1)$$

and by a natural extension of notation:

$$\text{COST}_{\mathcal{A}}^{(\omega, c)}(T) = \max_{x \in V} \text{COST}_{\mathcal{A}}^{(\omega, c)}(T, x).$$

The technical result which we will obtain in Section 4 may now be stated as follows.

Proposition 3.1. *For any $c \in \mathbb{N}$, $L \in \mathbb{N}$, there exists an algorithm running in time $(cn)^{O(L)}$, which for any tree T constructs a stable sequence assignment S and computes a value of ω such that $\omega \leq \frac{1}{L} \text{COST}_{\mathcal{A}_S}^{(\omega, c)}(T)$ and:*

$$\text{COST}_{\mathcal{A}_S}^{(\omega, c)}(T) \leq \left(1 + \frac{12}{c}\right) \text{OPT}(T).$$

In order to convert the obtained strategy \mathcal{A}_S with a small value of $\text{COST}^{(\omega, c)}$ into a strategy with small COST , we describe in Section 5 an appropriate strategy conversion mechanism. The approach we adopt is applicable to any strategy based on a stable sequence assignment and consists in concatenating, for each vertex $v \in V$, a prefix to the query sequence $S(v)$ in the form of a separately computed sequence $R(v)$, which does not depend on $S(v)$. The considered query sequences are thus of the form $R(v) \circ S(v)$, where the symbol ‘ \circ ’ represents sequence concatenation. Intuitively, the sequences R , taken over the whole tree, reflect the structure of a specific solution to the unweighted tree search problem on a contraction of tree T , in which each edge connecting a node to a child with weight at least $c\omega$ is contracted. We recall that the optimal number of queries to reach a target in an unweighted tree is $O(\log n)$, and the goal of this conversion is to reduce the number of light down queries in the combined strategy to at most $O(\log n)$.

Proposition 3.2. *For any fixed $\omega > 0$ there exists a polynomial-time algorithm which for a tree T computes a sequence assignment $R : V \rightarrow V^*$, such that, for any strategy \mathcal{A}_S based on a stable sequence assignment S , the sequence assignment S^+ , given by $S^+(v) = R(v) \circ S(v)$ for each $v \in V$, has the following property:*

$$\text{COST}_{\mathcal{A}_{S^+}}(T) \leq \text{COST}_{\mathcal{A}_S}^{(\omega, c)}(T) + 4(2c + 1)\omega \log_2 n.$$

The proof of Proposition 3.2 is provided in Section 5.

We are now ready to put together the two bounds. Combining the claims of Proposition 3.1 for $L = \lceil c^2 \log_2 n \rceil$ (with $\omega \leq \frac{1}{L} \text{COST}_{\mathcal{A}_S}^{(\omega, c)}(T) \leq \frac{\text{COST}_{\mathcal{A}_S}^{(\omega, c)}(T)}{c^2 \log_2 n}$) and Proposition 3.2, we obtain:

$$\begin{aligned} \text{COST}_{\mathcal{A}_{S^+}}(T) &\leq \text{COST}_{\mathcal{A}_S}^{(\omega, c)}(T) + 4(2c + 1)\omega \log_2 n \leq \text{COST}_{\mathcal{A}_S}^{(\omega, c)}(T) + 12c\omega \log_2 n \leq \\ &\leq \text{COST}_{\mathcal{A}_S}^{(\omega, c)}(T) + 12c \log_2 n \frac{\text{COST}_{\mathcal{A}_S}^{(\omega, c)}(T)}{c^2 \log_2 n} \leq \left(1 + \frac{12}{c}\right) \text{COST}_{\mathcal{A}_S}^{(\omega, c)}(T) \leq \\ &\leq \left(1 + \frac{12}{c}\right)^2 \text{OPT}(T) \leq \left(1 + \frac{168}{c}\right) \text{OPT}(T). \end{aligned}$$

After putting $\varepsilon = \frac{168}{c}$ and noting that in stating our result we can safely assume $c = O(\text{poly}(n))$ (beyond this, the tree search problem can be trivially solved optimally in $O(n^n)$ time using exhaustive search), we obtain the main theorem of the Section.

Theorem 3.3. *There exists an algorithm running in $n^{O\left(\frac{\log n}{\varepsilon^2}\right)}$ time, providing a $(1+\varepsilon)$ -approximation solution to the weighted tree search problem for any $0 < \varepsilon < 1$. \square*

3.2 Extension: A Poly-Time $O(\sqrt{\log n})$ -Approximation Algorithm

We now present the second main result of this work. By recursively applying the previously designed QPTAS (Theorem 3.3), we obtain a polynomial-time $O(\sqrt{\log n})$ -approximation algorithm for finding search strategy for an arbitrary weighted tree. We start by informally sketching the algorithm — we follow here the general outline of the idea from [7]. The algorithm is recursive and starts by finding a minimal subtree T^* of an input tree whose removal disconnects T into subtrees, each of size bounded by $n/2^{\sqrt{\log n}}$. The tree T^* will be processed by our optimal algorithm described in Section 3.1. This results either in locating the target node, if it belongs to T^* , or identifying the component of $T - T^*$ containing the target, in which case the search continues recursively in the component. However, for the final algorithm to have polynomial running time, the tree T^* needs to be of size $2^{O(\sqrt{\log n})}$. This is obtained by contracting paths in T^* (each vertex of the path has at most two neighbors in T^*) into single nodes having appropriately chosen weights. Since T^* has $2^{O(\sqrt{\log n})}$ leaves, this narrows down the size of T^* to the required level and we argue that an optimal search strategy for the ‘contracted’ T^* provides a search strategy for the original T^* that is within a constant factor from the cost of T^* .

A formal exposition and analysis of the obtained algorithm is provided in Section 6.

Theorem 3.4. *There is a $O(\sqrt{\log n})$ -approximation polynomial time algorithm for the weighted tree search problem.*

4 Proof of Proposition 3.1: Quasi-Polynomial Computation of Strategies with Small $\text{COST}^{(\omega, c)}$

4.1 Preprocessing: Time Alignment in Schedules

We adopt here a method similar but arguably more refined than rounding techniques in scheduling problems of combinatorial optimization, showing that we could discretise the starting and finishing time of jobs, as well as weights of vertices, in a way to restrict the size of state space for each node to $n^{O(\log n)}$, without introducing much error.

Fix $c \in \mathbb{N}$ and $\omega = \frac{a}{cn}$ for some $a \in \mathbb{N}$. (In subsequent considerations, we will have $c = \Theta(1/\varepsilon)$, $a = O(\frac{n}{\log n})$ and $\omega = \Omega(\varepsilon/\log n)$.) Given a tree $T = (V, E, w)$, let $T' = (V, E, w')$ be a tree with the same topology as T but with weights rounded up as follows:

$$w'(v) = \begin{cases} \lceil w(v) \rceil_{\omega}, & \text{if } w(v) > c\omega, \\ \lceil w(v) \rceil_{\frac{1}{cn}}, & \text{otherwise.} \end{cases} \quad (4.1)$$

We will informally refer to vertices with $w(v) > c\omega$ (equivalently $w'(v) > c\omega$) as *heavy vertices* and vertices with $w(v) \leq c\omega$ (equivalently $w'(v) \leq c\omega$) as *light vertices*. (Note that $w(v) \leq c\omega$ if and only if $w'(v) \leq c\omega$.) When designing schedules, we consider time divided into *boxes* of duration ω , with the i -th box equal to $[i\omega, (i+1)\omega]$. Each box is divided into a identical *slots* of length $\frac{1}{cn}$.

In the tree T' , the duration of a query to a heavy vertex is an integer number of boxes, and the duration of a query to a light vertex is an integer number of slots. We next show that, without affecting significantly the approximation ratio of the strategy, we can align each query to a heavy vertex in the schedule so that it occupies an interval of full adjacent boxes, and each query to a light vertex in the schedule so that it occupies an interval of full adjacent slots (possibly contained in more than one box).

We start by showing the relationship between the costs of optimal solutions for trees T and T' .

Lemma 4.1. $\text{OPT}(T) \leq \text{OPT}(T') \leq (1 + \frac{2}{c})\text{OPT}(T)$.

Proof. The inequality $\text{OPT}(T) \leq \text{OPT}(T')$ follows directly from the monotonicity of the cost of the solution with respect to vertex weights, since we have $w'(v) \geq w(v)$, for all $v \in V$.

To show the second inequality, we note that by the definition of weights (4.1), for any vertex v , $w'(v) \leq (1 + \frac{1}{c})w(v) + \frac{1}{cn}$.

Consider an optimal strategy \mathcal{O} for tree T and let $\mathbf{Q}_{\mathcal{O}}(T, x) = (v_1, \dots, v_k)$ be the time-ordering of queries performed by strategy \mathcal{O} on tree T to find a target vertex x . Let \mathcal{O}' be the strategy which follows the same time-ordering of queries when locating target x in T' . We have:

$$\begin{aligned} \text{COST}_{\mathcal{O}'}(T', x) &= \sum_{i=1}^k w'(v_i) \leq \sum_{i=1}^k \left(\left(1 + \frac{1}{c}\right) w(v) + \frac{1}{cn} \right) \leq \frac{1}{c} + \left(1 + \frac{1}{c}\right) \sum_{i=1}^k w(v) \leq \\ &\leq \left(1 + \frac{2}{c}\right) \text{OPT}(T), \end{aligned}$$

where we used the fact that, by Observation 2.1, $\text{OPT}(T) \geq 1$. Since $\text{OPT}(T') \leq \max_{x \in V} \text{COST}_{\mathcal{O}'}(T', x)$, the claim follows. \square

Lemma 4.2. *There exists a consistent schedule assignment \hat{S} for tree T' such that $\text{COST}_{\mathcal{A}_{\hat{S}}}(T') \leq (1 + \frac{3}{c})\text{OPT}(T')$ and for all $v \in V$ we have that*

- if $w'(v) > c\omega$, (v is heavy), then the starting time t of any job (v, t) in the schedule $\hat{S}(u)$ of any $u \in V$ is an integer multiple of ω (aligned to a box),
- if $w'(v) \leq c\omega$, (v is light), then the starting time t of any query (v, t) in the schedule $\hat{S}(u)$ of any $u \in V$ is an integer multiple of $\frac{1}{cn}$ (aligned to a slot).

Proof. We consider an optimal consistent schedule assignment $\hat{\Sigma}$ for tree T' , $|\hat{\Sigma}| = \text{OPT}(T')$. Fix $u \in V$ arbitrarily, and let $(v_{u,i}, t_{u,i})$ be the i -th query job in $\hat{\Sigma}(u)$. Consider now the schedule $\hat{\Sigma}^*(u)$ for T based on the same sequence assignment, in which the job $(v_{u,i}, t_{u,i})$ is replaced by the job $(v_{u,i}, t_{u,i}^*)$ with $t_{u,i}^* = (1 + \frac{2}{c})t_{u,i}$. We have for any two consecutive jobs at u :

$$t_{u,i+1}^* - t_{u,i}^* = \left(1 + \frac{2}{c}\right)(t_{u,i+1} - t_{u,i}) \geq \left(1 + \frac{2}{c}\right)w(v_{u,i}), \quad (4.2)$$

where we assume by convention that for the last job index i_{max} , $t_{u,i_{max}+1} = |\hat{\Sigma}(u)|$. We now observe that schedule assignment $\hat{\Sigma}^*$ on tree T can be directly converted into schedule assignment \hat{S} on tree T' as follows. The query sequence of each vertex is preserved unchanged. If $v_{u,i}$ is a heavy vertex, then within time interval $[t_{u,i}^*, t_{u,i+1}^*]$ we allocate to vertex $v_{u,i}$ an interval of full boxes, starting at time $\lceil t_{u,i}^* \rceil_\omega$. Indeed, by (4.2) we have:

$$t_{u,i+1}^* - \lceil t_{u,i}^* \rceil_\omega > t_{u,i+1}^* - t_{u,i}^* - \omega > \left(1 + \frac{2}{c}\right)w(v_{u,i}) - \omega > w(v_{u,i}) + \omega > w'(v_{u,i}).$$

Since no two jobs overlap and the time transformation is performed identically for all vertices, the validity and consistency of schedule assignment \hat{S} for tree T' follows. We also have $|\hat{S}| \leq (1 + \frac{2}{c})|\hat{\Sigma}| = (1 + \frac{2}{c})\text{OPT}(T')$.

To obtain the second part of the claim (alignment for light vertices) it suffices to round up the starting time of query times of all (light) vertices to an integer multiple of $\frac{1}{cn}$. Since all weights in T' are integer multiples of $\frac{1}{cn}$, and so are the starting times of queries to heavy vertices in \hat{S} , the correctness and consistency of the obtained schedule again follows directly. This final transformation increases the duration by at most $\frac{1}{c} \leq \frac{1}{c}\text{OPT}(T')$, and combining the bounds for both the transformations finally gives the claim. \square

A schedule on tree T' satisfying the conditions of Lemma 4.2, and the resulting search strategy, are called *aligned*. Subsequently, we will design an aligned strategy on tree T' , and compare the quality of the obtained solution to the best aligned strategy for T' .

The intuition between the separate treatment of heavy vertices (aligned to boxes) and light vertices (aligned to slots) in aligned schedules is the following. Whereas the time ordering of boxes is essential in the design of the correct strategy, in our dynamic programming approach we will not be concerned about the order of slots within a single box (i.e., the order of queries to light vertices placed in a single box). This allows us to reduce the state space of a node. Whereas the ordering of slots in the box will eventually have to be repaired to provide a correct strategy, this will not affect the quality of the overall solution too much (except for the issue of light down queries pointed out earlier, which are handled separately in Section 5).

4.2 Dynamic Programming Routine for Fixed Box Size

Let the values of parameter c and box size ω be fixed as before. Additionally, let $L \in \mathbb{N}$ be a parameter representing the time limit for the duration of the considered vertex schedules when measured in boxes, i.e., the longest schedule considered by the procedure will be of length $L\omega$ (we will eventually choose an appropriate value of $L = O(\log n)$ as required when showing Theorem 3.3).

Before presenting formally the considered quasi-polynomial time procedure, we start by outlining an (exponential time) algorithm which verifies if there exists an aligned schedule assignment $\hat{\Sigma}$ for T' whose duration is at most $L\omega$. Notice that since all weights in T' are integer multiples of $\frac{1}{cn}$, the optimal aligned schedule assignment will start and complete the execution of all queries at times which are integer multiples of $\frac{1}{cn}$; thus, we may restrict the considered class of schedules to those having this property. Any possible schedule of length at most $L\omega$ at a vertex v , which may appear in $\hat{\Sigma}$, will be represented in the form of the pair (σ_v, t_v) , where:

- σ_v is a Boolean array with $L\omega cn$ entries, where $\sigma_v[i] = 1$ when time slot $[\frac{i}{cn}, \frac{i+1}{cn}]$ is occupied in the schedule at v , and $\sigma_v[i] = 0$ otherwise.
- $t_v \in \mathbb{R}$ represents the start time of the query to v in the schedule of v (we put $t = \perp$ if such a query does not appear in the schedule).

We now state some necessary conditions for a consistent schedule, known from the analysis of the unweighted search problem (cf. e.g. [16, 30, 33]). The first observation expresses formally the constraint that the same time slot cannot be used in the schedules of two children of a node v , unless it is separated by an (earlier) query to node v itself. All time slots before the starting time t_v of job (v, t_v) are free if and only if the corresponding time slot is free for all of the children of v .

Observation 4.3. *Assume that the tuple $(\sigma_v, t_v)_{v \in V}$ corresponds to a consistent schedule. Let $v \in V$ be an arbitrarily chosen node with set of children $\{v_1, \dots, v_l\}$. Let the completion time t_{end}^v of the query to v in the schedule of v be given as:*

$$t_{end}^v = \begin{cases} t_v + w'(v), & \text{if } t_v \neq \perp, \\ +\infty, & \text{if } t_v = \perp. \end{cases}$$

Then, for any time slot $[\frac{i}{cn}, \frac{i+1}{cn}]$, we have:

$$\left. \begin{aligned} \sigma_v[i] &= \sum_{j=1}^l \sigma_{v_j}[i], & \text{when } \frac{i+1}{cn} \leq t_v, \\ \sigma_v[i] &= 1 \text{ and } \sum_{j=1}^l \sigma_{v_j}[i] = 0, & \text{when } t_v < \frac{i+1}{cn} \leq t_{end}^v, \\ \sigma_v[i] &= 0, & \text{when } \frac{i+1}{cn} > t_{end}^v. \end{aligned} \right\} \quad (4.3)$$

We remark that the last of the above conditions (4.3) follows from the w.l.o.g. assumption we made when defining sequence assignments that whenever node v appears in the schedule of v , it is the last node in the query sequence for v .

Moreover, any valid search strategy which locates a target vertex must eventually query at least one of the endpoints of every edge of the tree T' , since otherwise, it will not be able to distinguish targets located at these two endpoints. We thus make the following observation.

Observation 4.4. *Assume that the tuple $(\sigma_v, t_v)_{v \in V}$ represents a consistent schedule. Let $v \in V$ be an arbitrarily chosen node with set of children $\{v_1, \dots, v_l\}$. Then:*

$$\text{If } t_v = \perp, \text{ then } t_{v_j} \neq \perp, \text{ for all } 1 \leq j \leq l. \quad (4.4)$$

Conditions (4.3) and (4.4) provide us with necessary conditions which must be satisfied by any consistent aligned schedule assignment.

In order to lower-bound the duration of the consistent aligned schedule assignment with minimum cost, we perform an exhaustive bottom-up evaluation of aligned schedules which satisfy the constraints of (4.4), and a slightly weaker form of the constraints of (4.3). These weaker constraints are introduced to reduce the running time of the algorithm. Instead of considering individual slots of a schedule which may be empty or full, $\sigma_v[i] \in \{0, 1\}$, we consider the load of each box in the same schedule, defined as the proportion of occupied slots within the box. Formally, for the p -th box, $0 \leq p < L$, the *load* $s_v[p]$ is given as:

$$s_v[p] = \frac{1}{\omega cn} \sum_{i=p\omega cn}^{(p+1)\omega cn-1} \sigma_v[i], \quad s_v[p] \in \left\{0, \frac{1}{\omega cn}, \frac{2}{\omega cn}, \dots, 1\right\},$$

where we recall that ωcn is an integer by the choice of ω . We will call a box with load $s_v[p] = 0$ an *empty box*, a box with load $s_v[p] = 1$ a *full box*, and a box with load $0 < s_v[p] < 1$ a *partially full box* in the schedule of v .

By summing over all slots within each box, we obtain the following corollary directly from Observation 4.3.

Corollary 4.5. *Assume that the tuple $(s_v, t_v)_{v \in V}$ corresponds to a consistent schedule. Let $v \in V$ be an arbitrarily chosen node with set of children $\{v_1, \dots, v_l\}$ and completion time t_{end}^v of the query to v given as in Observation 4.3. Let a_p be the contribution to the load of the p -th box of the query job for vertex v , i.e.*

$$a_p = \begin{cases} \frac{1}{\omega} | [t_v, t_{end}^v] \cap [p\omega, (p+1)\omega] | & \text{if } t_v \neq \perp, \\ 0 & \text{if } t_v = \perp. \end{cases}$$

Then, for any box $[p\omega, (p+1)\omega]$, $0 \leq p < L$, we have:

$$\left. \begin{aligned} s_v[p] &= a_p + \sum_{j=1}^l s_{v_j}[p] \in [0, 1], & \text{when } t_{end}^v \geq (p+1)\omega, \\ s_v[p] &\geq a_p, & \text{when } p\omega < t_{end}^v < (p+1)\omega, \\ s_v[p] &= 0, & \text{when } t_{end}^v \leq p\omega. \end{aligned} \right\} \quad (4.5)$$

Moreover, for any box $[p\omega, (p+1)\omega]$, $0 \leq p < L$, we have:

$$\text{For all } 1 \leq j \leq l, \text{ the following bound holds: } s_{v_j}[p] + a_p \leq 1. \quad (4.6)$$

We remark that the statement of Corollary 4.5 treats specially one box, namely the one which contains strictly within it the time moment t_{end}^v . For this box, we are unable to make a precise statement about $s_v[p]$ based on the description of the schedules of its children, and content ourselves with a (potentially) weak lower bound $s_v[p] \geq a_p = \frac{1}{\omega}(t_{end}^v - p\omega)$. This is the direct reason for the slackness in our subsequent estimation, which loses ω time per down query. However, we note that by the definition of aligned schedule, a query to a heavy vertex will never begin or end strictly inside a box, and will not lead to the appearance of this issue. We remark that condition (4.6) additionally stipulates that within any box, it must be possible to schedule the contribution of the query to v and the contribution of any child v_j to the load of the box in a non-overlapping way. We now show that the shortest schedule assignments satisfying the set of constraints (4.4), (4.5), and (4.6) can be found in $n^{O(\log n)}$ time. This is

achieved by using the procedure BUILDSTRATEGY, presented in Algorithm 4.1, which returns for a node v a non-empty set of schedules $\hat{\mathcal{S}}[v]$, such that each $s_v \in \hat{\mathcal{S}}[v]$ can be extended into the sought assignment of schedules in its subtree, $(s_u, t_u)_{u \in V(T_v)}$. In the statement of Algorithm 4.1, we recall that, given a tree $T = (V, E, w)$, tree $T' = (V, E, w')$ is the tree with weights rounded up to the nearest multiple of the length of a slot (see Equation (4.1)).

The subsequent steps taken in procedure BUILDSTRATEGY can be informally sketched as follows. The input tree T' is processed in a bottom-up manner and hence, for an input vertex v , the recursive calls for its children v_1, \dots, v_l are first made, providing schedule assignments for the children (see lines 3-4). Then, the rest of the pseudocode is responsible for using these schedule assignments to obtain all valid schedule assignments for v . Lines 10-14 merge the schedules of the children in such a way that a set $\hat{\mathcal{S}}_i^*$, $i \in \{1, \dots, l\}$, contains all schedule assignments computed on the basis of the schedules for the children v_1, \dots, v_i . Thus, the set $\hat{\mathcal{S}}_l^*$ is the final product of this part of the procedure and is used in the remaining part. Note that a schedule assignment in $\hat{\mathcal{S}}_l^*$ may not be valid since a query to v is not accommodated in it — the rest of the pseudocode is responsible for taking each schedule $s \in \hat{\mathcal{S}}_l^*$ and inserting a query to v into s . More precisely, the subroutine INSERTVERTEX is used to place the query to v at all possible time points (depending whether v is heavy or light). We note that the subroutine MERGESCHEDULES, for each schedule s it produces, sets a Boolean ‘flag’ $s.\text{must_contain_}v$ that whenever equals *false*, indicates that querying v is not necessary in s to obtain a valid schedule for v (this happens if s queries all children of v). A detailed analysis of procedure BUILDSTRATEGY can be found in the proof of Lemma 4.6.

Lemma 4.6. *For fixed constants $L, c \in \mathbb{N}$, calling procedure BUILDSTRATEGY($r(T), \omega$), where $r(T)$ is the root of the tree, determines if there exists a tuple $(s_v, t_v)_{v \in V}$ which satisfies constraints (4.4), (4.5), and (4.6), or returns an empty set otherwise.*

Proof. The formulation of procedure BUILDSTRATEGY directly enforces that the constraints (4.4), (4.5), and (4.6) are fulfilled at each level of the tree, in a bottom-up manner.

For each vertex $v \in V$, we show by induction on the tree size that upon termination of procedure BUILDSTRATEGY(v, ω), the returned variable $\hat{\mathcal{S}}[v]$ is the set of all *minimal* schedules $(s_v, t_v) \in \hat{\mathcal{S}}[v]$ which can be extended within the subtree T_v to a data structure $(s_u, t_u)_{u \in V(T_v)}$, for some $(s_u, t_u) \in \hat{\mathcal{S}}[u]$, $u \in V(T_v)$, in such a way that the conditions (4.4), (4.5), and (4.6) hold within subtree T_v . Here, minimality of a schedule is a trivial technical assumption, understood in the sense of the following very restrictive partial order: we say $(s_v, t_v) \leq (s'_v, t'_v)$ if $s_v[p] \leq s'_v[p]$ for all $0 \leq p \leq L - 1$ and $t_v = t'_v$. (In the pseudocode, rather than write (s_v, t_v) as a pair variable, we include t_v within the structure s_v as its special field $s_v.t_v$.)

The algorithm proceeds to merge together exhaustively all possible choices of schedules $(s_{v_i}, t_{v_i}) \in \hat{\mathcal{S}}[v_i]$ of all children v_i of v , $1 \leq i \leq l$. The merge is performed by computing, for any fixed choice $(s_{v_i}, t_{v_i})_{1 \leq i \leq l}$, the combined load of each box in the resultant schedule s :

$$s[p] \leftarrow \sum_{i=1}^l s_{v_i}[p], \quad (4.7)$$

where, as a technicality, we also put $s[p] \leftarrow +\infty$ whenever we obtain excessive load in a box ($s[p] > 1$), as to avoid inflating the size of the state space and consequently, the running time of the algorithm. In Algorithm 4.1, the computation of $s[p]$ through the sum (4.7) proceeds by a processing of successive children v_i , $1 \leq i \leq l$, so that a schedule s stored in the data structure $\hat{\mathcal{S}}_i^*$ represents $s[p] = \sum_{j=1}^i s_{v_j}[p]$. The summation of load is performed within the subroutine MERGESCHEDULES, which merges a schedule $s_{orig} \in \hat{\mathcal{S}}_{i-1}^*$ with a schedule $s_{add} \in \hat{\mathcal{S}}[v_i]$ to obtain the new schedule $s \in \hat{\mathcal{S}}_i^*$.

Algorithm 4.1 Dynamic programming routine BUILDSTRATEGY for a tree T' . $L, c \in \mathbb{N}$ are global parameters. Subroutines MERGESCHEDULES and INSERTVERTEX are provided further on.

```

1: procedure BUILDSTRATEGY(vertex  $v$ , box size  $\omega \in \mathbb{R}$ )
2:    $l \leftarrow$  number of children of  $v$  in  $T'$  // Denote by  $v_1, \dots, v_l$  the children of  $v$ .
3:   for  $i = 1..l$  do
4:      $\hat{\mathcal{S}}[v_i] \leftarrow$  BUILDSTRATEGY( $v_i, \omega$ );
5:    $s \leftarrow 0^L$ 
6:    $s.max\_child\_load \leftarrow 0^L$ 
7:    $s.must\_contain\_v \leftarrow false$ 
8:    $\hat{\mathcal{S}}_0 \leftarrow \{s\}$  //  $\hat{\mathcal{S}}_0$  contains the schedule with no queries.
9:   // Inductively,  $\hat{\mathcal{S}}_i^*$  is based on merging schedules at  $v_1, \dots, v_i$ .
10:  for  $i = 1..l$  do
11:     $\hat{\mathcal{S}}_i^* \leftarrow \emptyset$ 
12:    for each schedule  $s \in \hat{\mathcal{S}}_{i-1}^*$  do
13:      for each schedule  $s_{add} \in \hat{\mathcal{S}}[v_i]$  do
14:         $\hat{\mathcal{S}}_i^* \leftarrow \hat{\mathcal{S}}_i^* \cup \text{MERGESCHEDULES}(s, s_{add}, \omega)$ ;
15:   $\hat{\mathcal{S}}[v] \leftarrow \emptyset$ 
16:  for each  $s \in \hat{\mathcal{S}}_l^*$  do
17:    if  $w'(v) > c\omega$  then //  $v$  is heavy
18:      for  $p = 0..L-1$  do //attempt to insert (into  $s$ ) query to  $v$  starting from time-box  $p$ 
19:         $\hat{\mathcal{S}}[v] \leftarrow \hat{\mathcal{S}}[v] \cup \text{INSERTVERTEX}(s, v, \omega, p \cdot \omega)$ 
20:    else //  $v$  is light
21:      for real  $t = 0..L \cdot \omega$  step  $\frac{1}{cn}$  do
22:        //attempt to insert (into  $s$ ) query to  $v$  at a slot from time  $t$ 
23:         $\hat{\mathcal{S}}[v] \leftarrow \hat{\mathcal{S}}[v] \cup \text{INSERTVERTEX}(s, v, \omega, t)$ 
24:    if  $s.must\_contain\_v = false$  then
25:       $\hat{\mathcal{S}}[v] \leftarrow \hat{\mathcal{S}}[v] \cup \text{INSERTVERTEX}(s, v, \omega, \perp)$ 
26:  return  $\hat{\mathcal{S}}[v]$ 

```

Eventually, the set of schedules $\hat{\mathcal{S}}_l^*$, obtained after merging the schedules of all children of v , contains an element s satisfying (4.7). Next, we test all possible values of $t_v \in \mathbb{R} \cup \{\perp\}$, which are feasible for an aligned schedule. These values depend on whether vertex v is heavy or light, for which t_v should represent the starting time of a box or slot, respectively. Using procedure INSERTVERTEX, we then set the load of each box following (4.5):

$$s_v[p] \leftarrow \begin{cases} a_p + \sum_{j=1}^l s_{v_j}[p], & \text{when } t_{end}^v \geq (p+1)\omega, \\ a_p, & \text{when } p\omega < t_{end}^v < (p+1)\omega, \\ 0, & \text{when } t_{end}^v \leq p\omega, \end{cases} \quad (4.8)$$

where a_p is defined as in (4.3). In the pseudocode of function INSERTVERTEX, for compactness we replace the second and third condition by equivalently setting $s_v[p] \leftarrow a_p$ when the first condition does not hold. We additionally constrain in procedures MERGESCHEDULES and INSERTVERTEX the possibility of the condition $t_v = \perp$ occurring by enforcing the constraints of (4.4) (corresponding of the setting of parameter $s.must_contain_v$ to $false$). Condition (4.6) is enforced through procedures MERGESCHEDULES and INSERTVERTEX using the auxiliary array $s.max_child_load[p]$, $0 \leq p \leq L-1$, defined so that $s.max_child_load[p] \leftarrow \max_{1 \leq j \leq l} s_{v_j}[p]$.

Since $\hat{\mathcal{S}}[v_i]$, for all $1 \leq i \leq l$, contains all minimal schedules satisfying (4.4), (4.5), and (4.6), the same holds for $\hat{\mathcal{S}}[v]$, which was constructed by enforcing only the required constraints. We remark that we obtain only the set of minimal (and not all) schedules due to the slight

Algorithm 4.2 Subroutines MERGESCHEDULES and INSERTVERTEX of procedure BUILDSTRATEGY from Algorithm 4.1.

```

1: procedure MERGESCHEDULES(schedule  $s_{orig}$ , schedule  $s_{add}$ , box size  $\omega \in \mathbb{R}$ )
2:    $s \leftarrow s_{orig}$  // copy schedule and its properties to answer
3:   for  $p = 0..L - 1$  do // for each time-box add load of  $s_1$  and  $s_2$ 
4:      $s[p] \leftarrow s_{orig}[p] + s_{add}[p]$ 
5:     if  $s[p] > 1$  then
6:        $s[p] \leftarrow +\infty$ 
7:      $s.max\_child\_load[p] \leftarrow \max\{s.max\_child\_load[p], s_{add}[p]\}$ 
8:     if  $s_{add}.t_v = \perp$  then
9:        $s.must\_contain\_v \leftarrow true$ 
10:  return  $s$ 
11:
12: procedure INSERTVERTEX(schedule  $s_{orig}$ , vertex  $v$ , box size  $\omega \in \mathbb{R}$ , time  $t \in \mathbb{R} \cup \{\perp\}$ )
13:   $s \leftarrow 0^L$  // initialize empty schedule for answer
14:  if  $t \neq \perp$  then
15:     $I \leftarrow [t, t + w'(v)]$  // time interval into which query to  $v$  is being inserted
16:     $s.t_v \leftarrow t$ 
17:     $t_{end}^v \leftarrow t + w'(v)$ 
18:  else
19:     $I \leftarrow \emptyset$ 
20:     $s.t_v \leftarrow \perp$ 
21:     $t_{end}^v \leftarrow +\infty$ 
22:  for  $p = 0..L - 1$  do // for each time-box
23:     $a_p \leftarrow \frac{1}{\omega} |I \cap [p \cdot \omega, (p + 1) \cdot \omega]|$  // contribution of query to  $v$  to load of box  $p$ 
24:    if  $s.max\_child\_load[p] + a_p > 1$  then
25:      return  $\emptyset$ 
26:    if  $t_{end}^v \geq (p + 1)\omega$  then
27:       $s[p] \leftarrow s_{orig}[p] + a_p$  // add load from children in box  $p$ 
28:      if  $s[p] > 1$  then //insertion failed
29:        return  $\emptyset$ 
30:    else
31:       $s[p] \leftarrow a_p$ 
32:  return  $\{s\}$ 

```

difference between (4.8) and (4.5) in the second condition: instead of requiring $s_v[p] \geq a_p$, we put $s_v[p] \leftarrow a_p$, thus setting the p -th coordinate of the schedule at its minimum possible value. \square

It follows directly from Lemma 4.6 that, for any value ω^* , tree T may only admit an aligned schedule assignment of duration at most ω^*L if a call to procedure BUILDSTRATEGY($r(T), \omega^*$) returns a non-empty set. Taking into account Lemmas 4.1 and 4.2, we directly obtain the following lower bound on the length of the shortest aligned schedule in tree T' .

Lemma 4.7. *If BUILDSTRATEGY($r(T), \omega^*$) = \emptyset , then:*

$$\omega^*L < \left(1 + \frac{3}{c}\right) \text{OPT}(T') \leq \left(1 + \frac{3}{c}\right) \left(1 + \frac{2}{c}\right) \text{OPT}(T) \leq \left(1 + \frac{11}{c}\right) \text{OPT}(T).$$

\square

Finally, we bound the running time of procedure BUILDSTRATEGY.

Lemma 4.8. *The running time of procedure $\text{BUILDSTRATEGY}(r(T), \omega)$ is at most $O((cn)^{\gamma L})$, for some absolute constant $\gamma = O(1)$, for any $\omega \leq n$.*

Proof. The procedure BUILDSTRATEGY is run recursively, and is executed once for each node of the tree. The time of each execution is upper-bounded, up to multiplicative factors polynomial in n , by the size of the largest of the schedule sets named $\hat{\mathcal{S}}[u]$, $u \in V$, or $\hat{\mathcal{S}}_i^*$, appearing in the procedure. We further focus only on bounding the size $|\hat{\mathcal{S}}|$ of the state space of distinct possible schedules in the (s_v, t_v) representation. The array s_v has size L , with each entry $s_v[p]$, $0 \leq p \leq L-1$, taking one of the values $s_v[p] \in \{0, \frac{1}{\omega cn}, \frac{2}{\omega cn}, \dots, 1\}$, where the size of the set of possible values is $\omega cn + 1 \in \mathbb{N}$. Additionally, in some of the auxiliary schedules, the additional array field $s_v.\text{max_child_load}$ has length L , with each entry $s_v.\text{max_child_load}[p]$, $0 \leq p \leq L-1$, likewise taking one of the values from the set $\{0, \frac{1}{\omega cn}, \frac{2}{\omega cn}, \dots, 1\}$. Finally, for the time t_v , we have: $t_v \in \{0, \omega, 2\omega, \dots, (L-1)\omega, \perp\}$, where the size of the set of possible values is $L+1$.

Overall, we obtain:

$$|\hat{\mathcal{S}}| \leq (L+1)(\omega cn + 1)^L (\omega cn + 1)^L \leq (L+1)(cn^2 + 1)^{2L} < (cn)^{L\gamma'},$$

where $\gamma' > 0$ is a suitably chosen absolute constant. Accommodating the earlier omitted multiplicative $O(\text{poly}(n))$ factors in the running time of the algorithm, we get the claim for some suitably chosen absolute constant $\gamma > \gamma'$. \square

4.3 Sequence Assignment Algorithm with Small $\text{COST}^{(\omega, c)}$

The procedure for computing a sequence assignment S which achieves a small value of $\text{COST}^{(\omega, c)}$ is given in Algorithm 4.3. This relies on procedure $\text{BUILDSTRATEGY}(r(T), \omega)$ as an essential subroutine, first determining the minimum value of $\omega = \frac{i}{cn}$, $i \in \mathbb{N}$, for which BUILDSTRATEGY produces a schedule. Since the schedule of a parent node v is based on an insertion of a query to v into the schedules of its children, a standard backtracking procedure allows us to determine the representation $(s_v, t_v)_{v \in V}$ of the schedules of all nodes of the tree.

Algorithm 4.3 Construction of sequence assignment S

```

1:  $\omega \leftarrow \frac{1}{cn}$ 
2: while  $\text{BUILDSTRATEGY}(r(T), \omega) = \emptyset$  do
3:    $\omega \leftarrow \omega + \frac{1}{cn}$ 
4:  $(s_v, t_v)_{v \in V} \leftarrow$  schedule assignment of duration at most  $L\omega$ , satisfying constraints (4.4), (4.5),
   and (4.6),
   reconstructed by backtracking through the sets  $(\hat{\mathcal{S}}[v])_{v \in V}$  computed in the last call
   to procedure  $\text{BUILDSTRATEGY}(r(T), \omega)$ .
5: for  $v \in V$  do
6:    $C(v) \leftarrow \emptyset$ 
7:   for  $u \in V(T_v)$  do
8:     if there is no vertex  $z \neq u$  on the path from  $v$  to  $u$  s.t.  $t_z < \lfloor t_u + w'(u) \rfloor_\omega + \omega$  then
9:        $C(v) \leftarrow C(v) \cup \{(\lfloor t_u \rfloor_\omega, \lfloor t_u + w'(u) \rfloor_\omega, u)\}$ 
10:   $S(v) \leftarrow$  sequence of vertices (third field) of  $C(v)$  sorted in non-decreasing
   order, with tuples compared by first field, then second field, then third field.
11: return  $(S(v))_{v \in V}$ 

```

We start by observing in Algorithm 4.3 that if a node v is not queried ($t_v = \perp$), then all of the children of v belong to the schedules produced by procedure BUILDSTRATEGY following condition (4.4), and thus they will also appear in $S(v)$. This guarantees the validity of the solution.

Lemma 4.9. *Algorithm 4.3 returns a correct query sequence assignment S for tree T .* \square

For the purposes of analysis, we extend the notion of backtracking procedure BUILDSTRATEGY in a natural way, so that, for every node $v \in V$ and box $0 \leq p \leq L - 1$, we describe precisely the contribution $c_v[p, u]$ of each vertex $u \in V(T_v)$ to the load $s_v[p]$. (See Fig. 4.1 for an illustration.)

Formally, for $u = v$ we have $c_v[p, v] \leftarrow a_p = |[t_v, t_v + w'(v)] \cap [p\omega, (p + 1)\omega]|$ if $t_v \neq \perp$, and $c_v[p, v] \leftarrow 0$, otherwise. Next, if $u \neq v$ and u belongs to the subtree of child v_i of v , we put:

$$c_v[p, u] \leftarrow \begin{cases} c_{v_i}[p, u], & \text{if } t_{end}^v > p\omega, \\ 0, & \text{otherwise,} \end{cases}$$

where the insertion time t_{end}^v for v is defined as in Observation 4.3. Comparing with (4.8), we have directly for all $0 \leq p < L$:

$$s_v[p] = \sum_{u \in V(T_v)} c_v[p, u].$$

Let $p_s(u)$ and $p_f(u)$ be the indices of the starting and final box, respectively, to which vertex u adds load, formally $p_s(u) = \min P_u$ and $p_f(u) = \max P_u$, where $P_u = \{p : |[t_u, t_u + w'(u)] \cap [p\omega, (p + 1)\omega]| > 0\}$. From the statement of Algorithm 4.3, we show immediately by inductive bottom-up argument that if $u \in S(v)$, then $\omega \sum_{p=p_s(u)}^{p_f(u)} c_v[p, u] = w'(u)$.

Lemma 4.10. *Let $(s_v, t_v)_{v \in T(V)}$ be a schedule assignment computed by BUILDSTRATEGY. For any vertices u and z such that (u, t_u) and (z, t_z) belong to the schedule at v , if either u or z is heavy, then $|[t_u, t_u + w'(u)] \cap [t_z, t_z + w'(z)]| = 0$.*

Proof. Note that procedure INSERTVERTEX is called for a heavy input vertex with its last parameter (insertion time) being a multiple of ω , and the weight $w'(u)$ is a multiple of ω by definition. Thus, the interval $[t_u, t_u + w'(u)]$ starts and ends at the beginning and end of a box, respectively. Hence, Constraint (4.6) gives the lemma. \square

As a consequence of Lemma 4.10, if these two jobs (u, t_u) and (z, t_z) overlap, where u and z belong to the sequence assignment $S(v)$, then both of the vertices u and z must be light, thus:

$$t_u > t_z - w'(u) - \omega = t_z + w'(z) - w'(z) - w'(u) - \omega \geq t_z + w'(z) - (2c + 1)\omega.$$


We now define the measure of progress $M(x, i)$ of strategy \mathcal{A}_S when searching for target x after i queries as follows. Let Q_i be the set of the first i queried vertices. Let v_i be the current root of the tree, $v_i = r(T \langle Q_i, x \rangle)$. Let $S_i(v) \subseteq S(v)$ be the subsequence (suffix) of $S(v)$ consisting of those vertices which have not yet been queried. Now, we define:

$$M(x, i) = \begin{cases} \min_{u \in S_i(v_i)} p_s(u), & \text{if } S_i(v_i) \neq \emptyset, \\ L, & \text{if } S_i(v_i) = \emptyset. \end{cases}$$

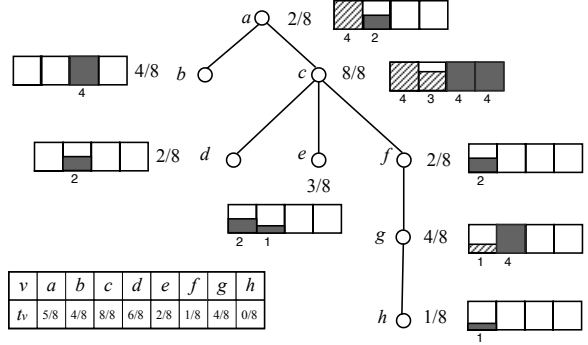
We have by definition, $M(x, i) \in \{0, 1, \dots, L - 1, L\}$. We obtain the next Lemma from a following straightforward analysis of the measure of progress: every time following sequence $S(v)$ we successively complete queries with an ‘up’ result with a total duration of at least a boxes, since the queried vertices are ordered in the first place according to minimum query time, and in the second place according to query duration, the value of the minimum $p_s(u)$, for $u \in S(v)$ remaining to be queried, advances by at least a boxes.

Lemma 4.11. *The measure of progress $M(x, i)$ has the following properties:*


1. *If the $(i + 1)$ -st query returns an ‘up’ result, then $M(x, i + 1) \geq M(x, i)$.*




(a) Tree T' with vertex weights.



(b) Sample schedule $(s_v, t_v)_{v \in V}$ obtained by backtracking procedure BuildStrategy with parameters $c = 1, n = 8, \omega = \frac{4}{8}$, (box size $\frac{4}{8}$, slot size $\frac{1}{8}$, 4 slots per box), $L = 4$. Note that the schedules $(s_v)_{v \in V}$ may correspond to different starting times of jobs within the prescribed box; the provided t_v are an example.



(c) Contribution of load of vertex e to different vertices of the tree.



(d) Sequences $S(v)$ computed by Algorithm 3 based on provided $(s_v, t_v)_{v \in V}$. Note that vertex e does not appear in $S(a)$ because of the query to c on the way.

Figure 4.1: Illustration of Algorithm 4.1 and 4.2. The depicted tree T' has vertex set $V = \{a, b, c, d, e, f, g, h\}$ and vertex weights:

v	a	b	c	d	e	f	g	h
$w'(v)$	$\frac{2}{8}$	$\frac{4}{8}$	$\frac{8}{8}$	$\frac{2}{8}$	$\frac{3}{8}$	$\frac{2}{8}$	$\frac{4}{8}$	$\frac{1}{8}$

2. If the $(i + 1)$ -st query returns a ‘down’ result, then $M(x, i + 1) \geq M(x, i) - (2c + 1)\omega$.
3. Suppose that between some two steps of the strategy, $i_2 > i_1$, each of the queries $(q_{i_1+1}, \dots, q_{i_2})$ returns an ‘up’ result, and moreover, the total cost of queries performed was at least $a\omega$, for some $a \in \mathbb{N}$:

$$\sum_{j=i_1+1}^{i_2} w'(q_j) \geq a\omega,$$

where $q_j = \mathbb{Q}_{\mathcal{A}_S, j}(T, x)$. Then, $M(x, i_2) \geq M(x, i_1) + a$.

□

Since the value of $M(x, i)$ is bounded from above by L , we obtain from Lemma 4.11 that the strategy \mathcal{A}_S necessarily terminates when looking for target x with cost at most $L\omega + (2c + 1)\omega d_x$,

$$\text{COST}_{\mathcal{A}_S}(T', x) \leq L\omega + (2c + 1)\omega d_x.$$

Thus, due to the definition of $\text{COST}^{(\omega, c)}$ in (3.1) and the monotonicity of the cost of a strategy with respect to vertex weights, we obtain the following:

Corollary 4.12. *For the sequence assignment computed by Algorithm 4.3 it holds*

$$\text{COST}_{\mathcal{A}_S}^{(\omega, c)}(T) \leq \text{COST}_{\mathcal{A}_S}^{(\omega, c)}(T') \leq \omega L.$$

□

To prove Proposition 3.1, it remains to show only the stability of the sequence assignment S .

Lemma 4.13. *The query sequence assignment S obtained by Algorithm 4.3 is stable.*

Proof. We perform the proof by induction. Following the definition of stability, assume that v is the root of the remaining subtree at some moment of executing \mathcal{A}_S on T' , and let u be a vertex such that u is a child of v lying on the path from v to the target x . We will show that following $S(u)$ always results in a subsequence of the sequence of queries performed by following $S(v)$.

Let $S^+(v)$ be the subsequence of vertices of $S(v)$ which lie in T_u , and let $S^-(v)$ be the subsequence of all remaining vertices of $S(v)$. Note that x belongs to T_u and hence any query to a node in $S^-(v)$ gives an ‘up’ reply.

We now observe the first (leftmost) difference v' of the sequences $S^+(v)$ and $S(u)$. Suppose that before such a difference occurs, the common fragment of the sequences contains a query to any vertex y on the path from u to x . Then, the root of both trees moves to the same child of y , and the process continues identically regardless of the initial root of the tree. Thus, such a vertex y cannot occur prior the difference in sequences $S^+(v)$ and $S(u)$.

Next, suppose that the first difference between the two sequences consists in the appearance of vertex v in sequence $S^+(v)$, i.e., $v' = v$. Then, the root of the tree moves from v to u , and the two processes proceed identically as required. This also implies that $t_v > t_u$.

Finally, we observe that no other first difference between the sequences $S^+(v)$ and $S(u)$ is possible by the formulation of Algorithm 4.3. In particular, if a triple $(\lfloor t_z \rfloor_\omega, \lceil t_z + w'(z) \rceil_\omega, z)$ is added to $C(u)$ in line 9, then the condition in line 8 and $t_v > t_u$ imply that the triple $(\lfloor t_z \rfloor_\omega, \lceil t_z + w'(z) \rceil_\omega, z)$ is added also to the set $C(v)$. Similarly, an insertion of a triple $(\lfloor t_z \rfloor_\omega, \lceil t_z + w'(z) \rceil_\omega, z)$ for $z \in V(T_u)$ into $C(v)$ implies that this triple also belongs to $C(u)$. Due to the sorting performed in line 10 of Algorithm 4.3, $S^+(v) = S(u)$.

The eventual deterministic coupling, which is obtained in all cases for the strategies starting at v and u , extends by induction to the execution of \mathcal{A}_S for trees rooted at a vertex v and its arbitrary descendant u' lying on the path from v to x , hence the claim holds. \square

For the chosen value ω , we can apply Lemma 4.7 with $\omega^* = \omega - \frac{1}{cn}$, obtaining:

$$\left(\omega - \frac{1}{cn}\right)L = \omega^*L < \left(1 + \frac{11}{c}\right)\text{OPT}(T),$$

thus, by Corollary 4.12,

$$\text{COST}_{\mathcal{A}_S}^{(\omega,c)}(T) \leq \left(1 + \frac{11}{c}\right)\text{OPT}(T) + \frac{L}{cn} \leq \left(1 + \frac{12}{c}\right)\text{OPT}(T),$$

where we took into account that trivially $L \leq n$ and $\text{OPT}(T) \geq 1$. We thus, by Lemmas 4.9, and 4.13 obtain the claim of Proposition 3.1.

5 Proof of Proposition 3.2: Reducing the Number of Down-Queries

We start with defining a function $\ell: V \rightarrow \{1, \dots, \lceil \log_2 n \rceil\}$ which in the following will be called a *labeling of T* and the value $\ell(v)$ is called the *label of v* . We say that a subset of nodes $H \subseteq V$ is an *extended heavy part* in T if $H = \{v\} \cup H'$, where all nodes in H' are heavy, no node in H' has a heavy neighbor in T that does not belong to H' and v is the parent of some node in H' . Let H_1, \dots, H_l be all extended heavy parts in T . Obtain a tree $T_C = (V_C, E_C)$ by contracting, in T , the subgraph H_i into a node denoted by h_i for each $i \in \{1, \dots, l\}$. In the tree T_C , we want to find its labeling $\ell': V_C \rightarrow \{1, \dots, \lceil \log_2 |V_C| \rceil\}$ that satisfies the following condition: for each two nodes u and v in V_C with $\ell'(u) = \ell'(v)$, the path between u and v has a node z satisfying $\ell'(z) < \ell'(u)$. One can obtain such a labeling by a following procedure that takes a subtree T'_C of T_C and an integer i as an input. Find a central node v in T'_C , set $\ell'(v) = i$ and call the procedure for each subtree T''_C of $T'_C - v$ with input T''_C and $i + 1$. The procedure is initially called for input T and $i = 1$. We also remark that, alternatively, such a labeling can be obtained via vertex rankings [16, 33].

Once the labeling ℓ' of T_C is constructed, we extend it to a labeling ℓ of T in such a way that for each node v of T we set $\ell(v) = \ell'(v)$ if $v \notin H_1 \cup \dots \cup H_l$ and $\ell(v) = \ell'(h_i)$ if $v \in H_i$, $i \in \{1, \dots, l\}$.

Having the labeling ℓ of T , we are ready to define a query sequence $R(v)$ for each node $v \in V$. The $R(v)$ contains all nodes u from T_v such that $\ell(u) < \ell(v)$ and each internal node z of the path connecting v and u in T satisfies $\ell(z) > \ell(u)$. Additionally, the nodes in $R(v)$ are ordered by increasing values of their labels. See Figure 5.1 for an example.

We start by making some simple observations regarding the sequence assignment R .

Observation 5.1. For each $v \in V$ and for each $u \in R(v)$, $w(u) \leq c\omega$. \square

Observation 5.2. For each $v \in V$, any two nodes in $R(v)$ have different labels. \square

Observation 5.3. The sequence assignment R can be computed in time $O(n \log n)$. \square

By x we refer to the target node in T . Fix S to be a stable sequence assignment in the remaining part of this section and by R we refer to the sequence assignment constructed above. Then, we fix S^+ to be $S^+(v) = R(v) \circ S(v)$ for each $v \in V$. Denote by U_i the first i nodes queried by \mathcal{A}_{S^+} and let $C_i = \min \ell(T \langle U_{i-1}, x \rangle)$ for each $i \geq 1$. For brevity we denote $U_0 = \emptyset$ and $C_0 = 0$; we also denote by u_i the node in $U_i \setminus U_{i-1}$, $i \geq 1$. A query made by \mathcal{A}_{S^+} to a node that belongs to $R(v)$ for some $v \in V$ is called an *R -query*; otherwise it is an *S -query*.

Lemma 5.4. For each $i \geq 0$, the nodes in $T\langle U_i, x \rangle$ with minimum label induce a connected subtree. \square

The next two lemmas will be used to conclude that the number of light queries performed by \mathcal{A}_{S^+} is bounded by $2 \log_2 n$ (see Lemma 5.7).

Lemma 5.5. If the i -th query of \mathcal{A}_{S^+} is an R -query resulting in an ‘up’ reply, then $C_{i+1} \geq C_i + 1$.

Proof. By construction, u_i has the minimum label among all nodes in $T\langle U_{i-1}, x \rangle$. By Lemma 5.4, either u_i is the unique node with label $\ell(u_i)$ in the tree $T\langle U_{i-1}, x \rangle$ or there are more nodes with this label and they all belong to a single extended heavy part in $T\langle U_{i-1}, x \rangle$ with u_i being closest to the root of $T\langle U_{i-1}, x \rangle$. In both cases, since the reply is ‘up’, we obtain that $T\langle U_i, x \rangle$ has no node with label $\ell(u)$, which proves the lemma. \square

Lemma 5.6. If the i -th query of \mathcal{A}_{S^+} is an R -query that results in a ‘down’ reply, then one of the two cases holds:

- (i) if u_i has no heavy child that belongs to $T\langle U_i, x \rangle$, then $C_{i+1} \geq C_i + 1$,
- (ii) if u_i has a heavy child that belongs to $T\langle U_i, x \rangle$, then $C_{i+1} = C_i = \ell(u_i)$ and for each $j > i$ such that $C_i = C_{j+1}$, all queries $i + 1, \dots, j$ are S -queries.

Proof. By Lemma 5.4, the nodes with label C_i induce a connected subtree in $T\langle U_{i-1}, x \rangle$. This immediately implies (i). By construction, if u_i has a heavy child u' that is in $T\langle U_i, x \rangle$, then $u' = r(T\langle U_i, x \rangle)$ and the labels of u_i and u' are the same. The latter is due to the fact that both u_i and u' belong to the same extended heavy part in T . Suppose for a contradiction that the j -th query (performed say on a node z) is an R -query and $C_{j+1} = C_i$, $j > i$. This in particular implies that $z \in R(r(T\langle U_{j-1}, x \rangle))$. Due to Lemma 5.5, the reply to this query is ‘down’. By (i), z has a heavy child that belongs to $T\langle U_j, x \rangle$. By Observation 5.1, z is a light node and therefore z along with some of its descendants and u_i with some of its descendants form two different extended heavy parts in T . Since z and u_i have the same label, there exists a light node u'_i in T on the path between u_i and z with label smaller than $\ell(u_i)$. Assume without loss of generality that no other node of this path that lies between u_i and u'_i has label smaller than $\ell(u'_i)$. The above-mentioned path is contained in $T\langle U_{i-1}, x \rangle$ since both u_i and z belong to this

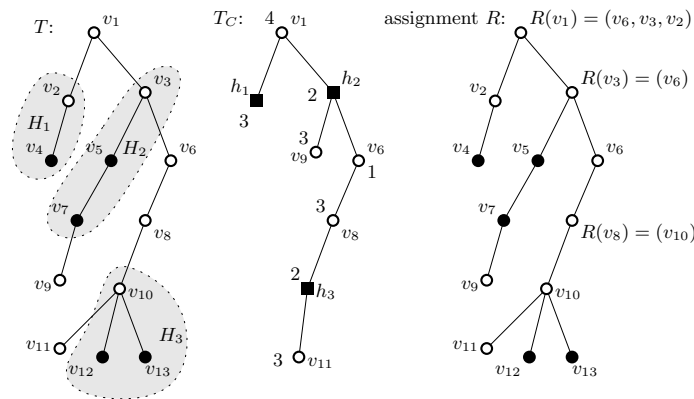


Figure 5.1: A tree T (on the left) has light vertices (marked as white nodes) and heavy ones (dark circles); also heavy extended parts are marked. The tree T_C (in the middle) is used together with its labeling (integers are the labels) to obtain the sequence assignment R (on the right); here we skip the sequence assignment for each node v for which $R(v) = \emptyset$.

subtree. This however implies that $u'_i \in R(v)$ because $\ell(u'_i) < \ell(u_i)$ and no node on the path between u_i and u'_i has label smaller than $\ell(u'_i)$. Moreover, u'_i precedes u_i in $R(v)$ meaning that among one for the first i queries, u'_i must have been queried — a contradiction with the fact that u'_i belongs to $T\langle U_i, x \rangle$. \square

Lemma 5.7. *For each target node, the total number of R -queries made by \mathcal{A}_{S^+} is at most $2 \log_2 n$.*

Proof. It follows from Lemmas 5.5 and 5.6 that after any two subsequent R -queries the value of parameter C_i increases by at least 1. \square

The next two lemmas will be used to bound the number of S -queries in S^+ receiving a ‘down’ reply to be at most $2 \log_2 n$.

Lemma 5.8. *If all nodes in $R(v)$ have been queried by \mathcal{A}_{S^+} after an i -th query for some $v \in V$ and v is the root of $T\langle U_i, x \rangle$, then $\ell(v) = C_{i+1}$.*

Proof. Suppose for a contradiction that $\ell(v) \neq C_{i+1}$. Since v belongs to $T\langle U_i, x \rangle$, we have that $\ell(v) > C_{i+1}$. Thus, by construction, there exists a light node u in $T\langle U_i, x \rangle$ with $\ell(u) = C_{i+1}$ such that all internal nodes on the path between v and u have labels larger than $\ell(u)$. Therefore, u belongs to $R(v)$ because v is the root of $T\langle U_i, x \rangle$. This implies that u has been already queried — a contradiction with u being in $T\langle U_i, x \rangle$. \square

Lemma 5.9. *If the i -th query of \mathcal{A}_{S^+} is an S -query performed on a light node and the reply is ‘down’, then $T\langle U_i, x \rangle$ has no light node with label C_i .*

Proof. Suppose that the i -th query is performed on a node u in $S(v)$ for some $v \in V$. Clearly, v is the root of $T\langle U_{i-1}, x \rangle$. Since the considered query is an S -query, all vertices in $R(v)$ have been already queried. Thus, by Lemma 5.8, $\ell(v) = C_i$. By construction, v is the only light node in this subtree having label C_i . Since the reply to the i -th query is ‘down’, v does not belong to $T\langle U_i, x \rangle$. \square

We are now ready to prove Proposition 3.2.

By Lemma 5.7, in \mathcal{A}_{S^+} , the total number of R -queries does not exceed $2 \log_2 n$. Note that since S is stable, for each target node x , the S -queries performed by \mathcal{A}_{S^+} are a subsequence of the queries performed by \mathcal{A}_S . Therefore, the potentially additional queries made by \mathcal{A}_{S^+} with respect to \mathcal{A}_S are R -queries. By Observation 5.1, each R -query is made on a light node. By definition of function $\text{COST}^{(\omega, c)}$ and Observation 5.1, any R -query increases the value of $\text{COST}^{(\omega, c)}$ of \mathcal{A}_{S^+} with respect to the value of $\text{COST}^{(\omega, c)}$ of \mathcal{A}_S by at most $(2c + 1)\omega$. Hence we have:

$$\text{COST}_{\mathcal{A}_{S^+}}^{(\omega, c)}(T) \leq \text{COST}_{\mathcal{A}_S}^{(\omega, c)}(T) + 2(2c + 1)\omega \log_2 n.$$

By Lemmas 5.6 and 5.9, the total number of queries in strategy \mathcal{A}_{S^+} to light nodes receiving ‘down’ replies can be likewise bounded by $2 \log_2 n$. Since each such query introduces a rounding difference of at most $(2c+1)\omega$ when comparing cost functions COST and $\text{COST}^{(\omega, c)}$, we thus obtain:

$$\text{COST}_{\mathcal{A}_{S^+}}(T) \leq \text{COST}_{\mathcal{A}_{S^+}}^{(\omega, c)}(T) + 2(2c + 1)\omega \log_2 n.$$

Combining the above observations gives the claim of the Proposition.

6 Proof of Theorem 3.4: A $O(\sqrt{\log n})$ -Approximation Algorithm

We start with some notation. Given a tree $T = (V, E, w)$ and a fixed value of parameter α , we find a subtree $T^* = (V^*, E^*)$ of the input tree T , called an α -separating tree, that satisfies: $r(T^*) = r(T)$ and each connected component of $T \setminus V^*$ has at most α vertices. An α -separating tree T^* is *minimal* if the removal of any leaf from T^* gives an induced tree that is not an α -separating tree. Then, for a target node $x \in V$, we introduce a recursive strategy \mathcal{R} that takes the following steps:

1. \mathcal{R} first applies strategy \mathcal{A}^* restricted to tree T^* to locate the node x' of T^* which is closest to the target x .
2. Then, \mathcal{R} queries x' , which either completes the search in case when x' is the target or provides a neighbor x'' of x' that is closer to the target than x' .
3. If x' is not the target, then the strategy calls itself recursively on the subtree $T_{x''}$ of $T \setminus \{x'\}$ containing x . The latter strategy for $T_{x''}$ is denoted by $\mathcal{R}_{x''}$. (Note that $T_{x''}$ is a connected component in $T \setminus V^*$.)

Such a search strategy \mathcal{R} obtained from \mathcal{A}^* and strategies $\mathcal{R}_{r(T')}$ (constructed recursively) for subtrees T' in $T \setminus V^*$ is called a $(\mathcal{A}^*, \{\mathcal{R}_{r(T')} \mid T' \in \mathcal{C}(T \setminus V^*)\})$ -strategy, where $\mathcal{C}(T \setminus V^*)$ is the set of connected components (subtrees) in $T \setminus V^*$.

The following bound on the cost of the strategy \mathcal{R} follows directly from the construction:

Lemma 6.1. *For a $(\mathcal{A}^*, \{\mathcal{R}_{r(T')} \mid T' \in \mathcal{C}(T \setminus V^*)\})$ -strategy \mathcal{R} for T it holds*

$$\text{COST}_{\mathcal{R}}(T) \leq \text{COST}_{\mathcal{A}^*}(T^*) + \max_{x' \in V^*} w(x') + \max_{T' \in \mathcal{C}(T \setminus V^*)} \text{COST}_{\mathcal{R}_{r(T')}}(T').$$

□

We now formally describe and analyze the aforementioned contractions of subpaths in a tree. A maximal path with more than one node in a tree T that consists only of vertices that have degree two in T is called a *long chain in T* . For each long chain P , contract it into a single node v_P with weight $\min_{u \in V(P)} w(u)$, obtaining a tree $\xi(T)$. In what follows, the tree $\xi(T)$ is called a *chain-contraction of T* .

Our first step is a remark that, at the cost of losing a multiplicative constant in the final approximation ratio, we may restrict ourselves to trees that have no long chains. This is due to the following observation.

Lemma 6.2. *Let T be a tree. Given a p -approximate search strategy for $\xi(T)$, a $(p + 1)$ -approximate search strategy for T can be computed in polynomial time.*

Proof. Let \mathcal{A}' be a search strategy for $\xi(T)$. We obtain a search strategy \mathcal{A} for T in two stages. In the first stage we ‘mimic’ the behavior of \mathcal{A}' : (i) if \mathcal{A}' queries a node v that also belongs to T , then \mathcal{A} also queries v ; (ii) if \mathcal{A}' queries a node v_P that corresponds to some long chain P in T , then \mathcal{A} queries, in T , a node with minimum weight in P . Note that after the first stage, the search strategy either located the target or determined that the target belongs to a subpath P' of some long chain P of T . Moreover, the total cost of all queries performed in the first stage is at most $\text{COST}_{\mathcal{A}'}(\xi(T))$.

Then, in the second stage we compute (in $O(n^2)$ -time) an optimal search strategy $\mathcal{A}_{P'}$ for P' [6]. Due to the monotonicity of the cost over taking subgraphs, $\text{COST}_{\mathcal{A}_{P'}}(P') = \text{OPT}(P') \leq \text{OPT}(T)$.

Both stages provide us with a search strategy for T with cost at most $\text{COST}_{\mathcal{A}'}(\xi(T)) + \text{OPT}(T)$. Since, $\text{OPT}(\xi(T)) \leq \text{OPT}(T)$ and $\text{COST}_{\mathcal{A}'}(\xi(T)) \leq p \cdot \text{OPT}(\xi(T))$, the lemma follows. □

Note that it is straightforward to verify whether any vertex v of T is a leaf in the α -separating tree of T and hence we obtain the following.

Observation 6.3. *Given a tree T with no long chain and α , a minimal α -separating tree of T can be computed in polynomial-time. \square*

Using Lemma 6.2 and choosing appropriately the value of α , one can obtain an α -separating tree of T having at most $t = 2^{O(\sqrt{\log n})}$ vertices.

Lemma 6.4. *Let T be any tree and let α be selected arbitrarily. If T^* is a minimal α -separating tree of T , then $\xi(T^*)$ has at most $4 \lceil \frac{n}{\alpha} \rceil$ vertices.*

Proof. By definition, for each leaf v of T^* , the subtree T_v has more than α nodes. Since these trees are node-disjoint, we obtain that there are at most $\lceil \frac{n}{\alpha} \rceil$ leaves in T^* . We denote the leaves of T^* by v_1, v_2, \dots, v_l , $l \leq \lceil \frac{n}{\alpha} \rceil$; note that $\xi(T^*)$ has the same leaves as T^* .

Let $V(\xi(T^*))$ be the vertex set of $\xi(T^*)$. Then, we claim that $|V(\xi(T^*))| = O(\lceil \frac{n}{\alpha} \rceil)$ by counting the number of nodes with different degrees in $\xi(T^*)$. Clearly, we have $|\{v \in V(\xi(T^*)) \mid \deg(v) > 2\}| \leq \lceil \frac{n}{\alpha} \rceil$. Since the tree $\xi(T^*)$ contains no long chains, the parent (if exists) of every node with degree exactly 2 must have degree at least 3. Thus,

$$|\{v \in V(\xi(T^*)) \mid \deg(v) = 2\}| \leq |\{v \in V(\xi(T^*)) \mid \deg(v) > 2\}| + 1 \leq \lceil \frac{n}{\alpha} \rceil + 1.$$

Hence we get $|V(\xi(T^*))| \leq 4 \lceil \frac{n}{\alpha} \rceil$. \square

With Lemmas 6.2, 6.4 and Observation 6.3 we are now ready to obtain the efficient recursive decomposition of the problem:

Lemma 6.5. *If there is a $O(1)$ -approximation algorithm running in $n^{O(\log n)}$ time for any input tree, then one can obtain a $O(\sqrt{\log n})$ -approximation algorithm with polynomial running time for any input tree.*

Proof. Suppose SOLVE is a given constant-factor approximation algorithm running in time $n^{O(\log n)}$ that, for any input tree T , outputs a search strategy for T . We then design a polynomial-time procedure REC as shown in Algorithm 6.1, which outputs a search strategy \mathcal{R} for an input tree T .

Algorithm 6.1 $O(\sqrt{\log n})$ -approximation procedure REC based on $n^{O(\log n)}$ -time constant approximation algorithm SOLVE

```

1: procedure REC(tree  $T = (V, E, w)$ )
2:    $n \leftarrow |V|$ 
3:   if  $n \leq 2^{\sqrt{\log n}}$  then
4:     return SOLVE( $T$ )
5:   else
6:      $\alpha \leftarrow n/2^{\sqrt{\log n}}$ 
7:      $T^* \leftarrow$  a minimal  $\alpha$ -separating tree of  $T$  with vertex set  $V^*$ 
8:      $\mathcal{A}^* \leftarrow$  SOLVE( $\xi(T^*)$ )
9:      $\mathcal{A}_{T^*} \leftarrow$  search strategy for  $T^*$  obtained from  $\mathcal{A}^*$  as described in proof of Lemma 6.2
10:    for each  $T'$  in  $\mathcal{C}(T \setminus V^*)$  do
11:       $\mathcal{R}_{r(T')} \leftarrow$  REC( $T'$ );
12:    return  $(\mathcal{A}_{T^*}, \{\mathcal{R}_{r(T')} \mid T' \in \mathcal{C}(T \setminus V^*)\})$ -strategy for  $T$ 

```

Each call to SOLVE in line 4 has running time $(2^{\sqrt{\log n}})^{O(\log(2^{\sqrt{\log n}}))}$, which is a polynomial in n . The same holds for each call in line 8 because, by Lemma 6.4, $\xi(T^*)$ has at most $4 \lceil \frac{n}{\alpha} \rceil = O(2^{\sqrt{\log n}})$ vertices. Thus, procedure REC has polynomial running time and it remains to bound the cost of the search strategy \mathcal{R} computed by REC.

To bound the recursion depth of REC, note that each time a recursive call is made, the size of instance (input tree) decreases $2^{\sqrt{\log n}}$ times. Thus, the depth is bounded by $\log_{(2^{\sqrt{\log n}})} n = \sqrt{\log n}$. In the search strategy computed by procedure REC, at each level of the recursion we execute the search strategy computed by one call to SOLVE and one vertex of the $(n/2^{\sqrt{n}})$ -separating tree is queried. This follows from the definition of $(\mathcal{A}_{T^*}, \{\mathcal{R}_{r(T')} \mid T' \in \mathcal{C}(T \setminus V^*)\})$ -strategy. By Lemma 6.2,

$$\text{COST}_{\mathcal{A}_{T^*}}(T^*) \leq c' \cdot \text{OPT}(T^*)$$

for some constant c' . By Lemma 6.1 and since $\text{OPT}(T^*) \leq \text{OPT}(T)$, the cost of \mathcal{R} at each recursion level is bounded by $(c'+1)\text{OPT}(T)$. This gives that $\text{COST}_{\mathcal{R}}(T) \leq c'\sqrt{\log n} \cdot \text{OPT}(T)$ as required. \square

Noting that the existence of a constant-approximation procedure with $n^{O(\log n)}$ running time follows from Theorem 3.3 (by taking $\varepsilon = 1$), the claim of Theorem 3.4 follows directly from Lemma 6.5.

Acknowledgment

The authors thank Jakub Łacki for preliminary discussions on the studied problem.

References

- [1] Esther M. Arkin, Henk Meijer, Joseph S. B. Mitchell, David Rappaport, and Steven Skiena. Decision trees for geometric models. *Int. J. Comput. Geometry Appl.*, 8(3):343–364, 1998.
- [2] Yosi Ben-Asher and Eitan Farchi. The cost of searching in general trees versus complete binary trees. Technical report, Technical report, 1997.
- [3] Michael Ben-Or and Avinatan Hassidim. The bayesian learner is optimal for noisy binary search (and pretty good for quantum as well). In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 221–230, 2008.
- [4] Renato Carmo, Jair Donadelli, Yoshiharu Kohayakawa, and Eduardo Sany Laber. Searching in random partially ordered sets. *Theor. Comput. Sci.*, 321(1):41–57, 2004.
- [5] Ferdinando Cicalese, Tobias Jacobs, Eduardo Sany Laber, and Marco Molinaro. On the complexity of searching in trees and partially ordered structures. *Theor. Comput. Sci.*, 412(50):6879–6896, 2011.
- [6] Ferdinando Cicalese, Tobias Jacobs, Eduardo Sany Laber, and Caio Dias Valentim. The binary identification problem for weighted trees. *Theor. Comput. Sci.*, 459:100–112, 2012.
- [7] Ferdinando Cicalese, Balázs Keszegh, Bernard Lidický, Dömötör Pálvölgyi, and Tomás Valla. On the tree search problem with non-uniform costs. *CoRR*, abs/1404.4504, 2014.
- [8] Dariusz Dereniowski. Edge ranking of weighted trees. *Discrete Applied Mathematics*, 154(8):1198–1209, 2006.

- [9] Dariusz Dereniowski. Edge ranking and searching in partial orders. *Discrete Applied Mathematics*, 156(13):2493–2500, 2008.
- [10] Dariusz Dereniowski and Marek Kubale. Efficient parallel query processing by graph ranking. *Fundam. Inform.*, 69(3):273–285, 2006.
- [11] Dariusz Dereniowski and Adam Nadolski. Vertex rankings of chordal graphs and weighted trees. *Inf. Process. Lett.*, 98(3):96–100, 2006.
- [12] Ehsan Emamjomeh-Zadeh, David Kempe, and Vikrant Singhal. Deterministic and probabilistic binary search in graphs. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 519–532, 2016.
- [13] Uriel Feige, Prabhakar Raghavan, David Peleg, and Eli Upfal. Computing with noisy information. *SIAM J. Comput.*, 23(5):1001–1018, 1994.
- [14] Archontia C. Giannopoulou, Paul Hunter, and Dimitrios M. Thilikos. Lifo-search: A min-max theorem and a searching game for cycle-rank and tree-depth. *Discrete Applied Mathematics*, 160(15):2089–2097, 2012.
- [15] Brent Heeringa, Marius Catalin Iordan, and Louis Theran. Searching in dynamic tree-like partial orders. In *Algorithms and Data Structures - 12th International Symposium, WADS 2011, New York, NY, USA, August 15-17, 2011. Proceedings*, pages 512–523, 2011.
- [16] Ananth V. Iyer, H. Donald Ratliff, and Gopalakrishnan Vijayan. Optimal node ranking of trees. *Inf. Process. Lett.*, 28(5):225–229, 1988.
- [17] Ananth V. Iyer, H. Donald Ratliff, and Gopalakrishnan Vijayan. Parallel assembly of modular products – an analysis. Technical report, Technical Report 88-86, Georgia Institute of Technology, 1988.
- [18] Richard M. Karp and Robert Kleinberg. Noisy binary search and its applications. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana, USA, January 7-9, 2007*, pages 881–890, 2007.
- [19] Meir Katchalski, William McCuaig, and Suzanne M. Seager. Ordered colourings. *Discrete Mathematics*, 142(1-3):141–154, 1995.
- [20] Eduardo Sany Laber, Ruy Luiz Milidiú, and Artur Alves Pessoa. On binary searching with nonuniform costs. *SIAM J. Comput.*, 31(4):1022–1047, 2002.
- [21] Eduardo Sany Laber and Marco Molinaro. An approximation algorithm for binary searching in trees. *Algorithmica*, 59(4):601–620, 2011.
- [22] Eduardo Sany Laber and Loana Tito Nogueira. Fast searching in trees. *Electronic Notes in Discrete Mathematics*, 7:90–93, 2001.
- [23] Eduardo Sany Laber and Loana Tito Nogueira. On the hardness of the minimum height decision tree problem. *Discrete Applied Mathematics*, 144(1-2):209–212, 2004.
- [24] Tak Wah Lam and Fung Ling Yue. Optimal edge ranking of trees in linear time. *Algorithmica*, 30(1):12–33, 2001.

- [25] Nathan Linial and Michael E. Saks. Searching ordered structures. *J. Algorithms*, 6(1):86–103, 1985.
- [26] Joseph W. H. Liu. Computational models and task scheduling for parallel sparse cholesky factorization. *Parallel Computing*, 3(4):327–342, 1986.
- [27] Joseph W. H. Liu. The role of elimination trees in sparse factorization. *SIAM. J. Matrix Anal. & Appl.*, 11(1):134–172, 1990.
- [28] Shay Mozes, Krzysztof Onak, and Oren Weimann. Finding an optimal tree searching strategy in linear time. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, pages 1096–1105, 2008.
- [29] Jaroslav Nešetřil and Patrice Ossona de Mendez. Tree-depth, subgraph coloring and homomorphism bounds. *Eur. J. Comb.*, 27(6):1022–1041, 2006.
- [30] Krzysztof Onak and Pawel Parys. Generalization of binary search: Searching in trees and forest-like partial orders. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21-24 October 2006, Berkeley, California, USA, Proceedings*, pages 379–388, 2006.
- [31] Alex Pothén. The complexity of optimal elimination trees. Technical report, Technical Report CS-88-13, Pennsylvania State University, 1988.
- [32] Ronald L. Rivest, Albert R. Meyer, Daniel J. Kleitman, Karl Winklmann, and Joel Spencer. Coping with errors in binary search procedures. *J. Comput. Syst. Sci.*, 20(3):396–404, 1980.
- [33] Alejandro A. Schäffer. Optimal node ranking of trees in linear time. *Inf. Process. Lett.*, 33(2):91–96, 1989.
- [34] George Steiner. Searching in 2-dimensional partial orders. *J. Algorithms*, 8(1):95–105, 1987.
- [35] Jayme Luiz Szwarcfiter, Gonzalo Navarro, Ricardo A. Baeza-Yates, Joísa de S. Oliveira, Walter Cunto, and Nivio Ziviani. Optimal binary search trees with costs depending on the access paths. *Theor. Comput. Sci.*, 290(3):1799–1814, 2003.