



New Partitioning Techniques and Faster Algorithms for Approximate Interval Scheduling

Spencer Compton¹ · Slobodan Mitrović² · Ronitt Rubinfeld³

Received: 18 August 2023 / Accepted: 12 July 2024 / Published online: 18 July 2024
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

Abstract

Interval scheduling is a basic algorithmic problem and a classical task in combinatorial optimization. We develop techniques for partitioning and grouping jobs based on their starting/ending times, enabling us to view an instance of interval scheduling on *many* jobs as a union of multiple interval scheduling instances, each containing only *a few* jobs. Instantiating these techniques in a dynamic setting produces several new results. For $(1 + \varepsilon)$ -approximation of job scheduling of n jobs on a single machine, we develop a fully dynamic algorithm with $O(\log n/\varepsilon)$ update and $O(\log n)$ query worst-case time. Our techniques are also applicable in a setting where jobs have weights. We design a fully dynamic *deterministic* algorithm whose worst-case update and query times are $\text{poly}(\log n, \frac{1}{\varepsilon})$. This is the *first* algorithm that maintains a $(1 + \varepsilon)$ -approximation of the maximum independent set of a collection of weighted intervals in $\text{poly}(\log n, \frac{1}{\varepsilon})$ time updates/queries. This is an exponential improvement in $1/\varepsilon$ over the running time of

We thank Benjamin Qi (MIT) for helpful discussions. S. Compton was supported in part by the National Defense Science & Engineering Graduate (NDSEG) Fellowship Program. S. Mitrović was supported by the Swiss National Science Foundation grant No. P400P2_191122/1, National Science Foundation award CCF-1733808, and FinTech@CSAIL. R. Rubinfeld was supported by the National Science Foundation TRIPDS program (Awards CCF-1740751 and DMS-2022448), National Science Foundation award CCF-2006664, and FinTech@CSAIL. This work was originally presented at the 50th International Colloquium on Automata, Languages, and Programming (ICALP 2023) conference.

-
- ✉ Spencer Compton
comptons@stanford.edu
 - ✉ Slobodan Mitrović
smitrovic@ucdavis.edu
 - ✉ Ronitt Rubinfeld
ronitt@csail.mit.edu

- 1 Computer Science Department, Stanford University, 353 Jane Stanford Way, Stanford, CA 94305, USA
- 2 Computer Science Department, UC Davis, Kemper Hall, Davis, CA 95616, USA
- 3 CSAIL, MIT, 32 Vassar Street, Cambridge, MA 02139, USA

an algorithm of Henzinger, Neumann, and Wiese [SoCG, 2020]. Our approach also removes all dependence on the values of the jobs' starting/ending times and weights.

Keywords Interval scheduling · Job scheduling · Dynamic algorithms · Local computation algorithms

1 Introduction

Job scheduling is a fundamental task in optimization, with applications ranging from resource management in computing [1, 2] to operating transportation systems [3]. Given a collection of *machines* and a set of *jobs* (or tasks) to be processed, the goal of job scheduling is to assign those jobs to the machines while respecting certain constraints. Constraints set on jobs may significantly vary. In some cases, a job has to be scheduled, but the starting time of its processing is not pre-specified. In other scenarios, a job can only be scheduled at a given time, but there is flexibility on whether to process the job or not. Frequent objectives for this task can include either maximizing the number of scheduled jobs or minimizing the needed time to process all given jobs.

An important variant of job scheduling is the task of *interval scheduling*: here, each job has a specified starting time and length, but a job is not required to be scheduled. Given M machines, the goal is to schedule as many jobs as possible. More generally, each job is also assigned a *reward* or weight, which can be thought of as a payment received for processing the given job. If a job is not processed, the payment is zero, i.e., there is no penalty. We refer to this variant as *weighted interval scheduling*. This problem, in a natural way, captures real-life scenarios. For instance, consider assigning crew members to flights where we aim to assign crews for as many flights as possible. In the context of interval scheduling, flights can be seen as jobs and the crew members as machines [3, 4]. Interval scheduling also has applications in geometrical tasks – it can be seen as a task of finding a collection of non-overlapping geometric objects. In this context, its prominent applications are in VLSI design [5] and map labeling [6, 7].

The aforementioned scenarios are executed in different computational settings. For instance, some use cases are dynamic in nature, e.g., a flight gets canceled. Then, in certain cases, we have to make online decisions, e.g., a customer must know immediately whether we are able to accept its request or not. In some applications, there might be so many requests that we would like to design extremely fast ways of deciding whether a given request/job can be scheduled, e.g., providing an immediate response to a user submitting a job for execution in a cloud.

1.1 The Computation Model

In our work, we focus on the dynamic setting of computation. Our algorithms for the fully dynamic setting design data structures that maintain an approximately optimal solution to an instance of the interval scheduling problem while supporting insertions

and deletions of jobs/intervals. The data structures also support queries of the maintained solution’s total weight and whether or not a particular interval is used in the maintained solution.

1.2 Our Results

Our first result, given in Sect. 4, focuses on designing an efficient dynamic algorithm for unweighted interval scheduling on a single machine. Prior to our work, the state-of-the-art result for this unweighted interval scheduling problem was due to [8], who design an algorithm with $O(\log n/\varepsilon^2)$ update and query time. We provide an improvement in the dependence on ε .

Theorem 1.1 (Unweighted dynamic, single machine) *Let \mathcal{J} be a set of n jobs. For any $\varepsilon > 0$, there exists a fully dynamic algorithm for $(1 + \varepsilon)$ -approximate unweighted interval scheduling for \mathcal{J} on a single machine performing updates in $O\left(\frac{\log(n)}{\varepsilon}\right)$ and queries in $O(\log(n))$ worst-case time.*

Theorem 1.1 can be seen as a warm-up for our most challenging and technically involved result, which is an algorithm for the dynamic *weighted* interval scheduling problem on a single machine. We present our approach in detail in Sect. 5. As a function of $1/\varepsilon$, our result is an exponential improvement compared to the running times obtained in [9]. We also remove all dependence on the job starting/ending times (previous work crucially used assumptions on the coordinates to bound the ratio of jobs’ lengths by a parameter N), and remove all dependence on the value of the job rewards.

Theorem 1.2 (Weighted dynamic, single machine) *Let \mathcal{J} be a set of n weighted jobs. For any $\varepsilon > 0$, there exists a fully dynamic algorithm for $(1 + \varepsilon)$ -approximate weighted interval scheduling for \mathcal{J} on a single machine performing updates and queries in worst-case time $T \in \text{poly}(\log n, \frac{1}{\varepsilon})$. The exact complexity of T is given by*

$$O\left(\frac{\log^{12}(n)}{\varepsilon^7} + \frac{\log^{13}(n)}{\varepsilon^6}\right).$$

1.3 Related Work

The closest prior work to ours is that of Henzinger et al. [9], and Bhore et al. [8]. The work of Henzinger et al. studies $(1 + \varepsilon)$ -approximate dynamic interval scheduling for one machine in both the weighted and unweighted setting. Unlike our main result in Theorem 1.2, they assume that: jobs have rewards within $[1, W]$; jobs have length at least 1; and jobs start/end within times $[0, N]$. They obtain deterministic algorithms with $O(\exp(1/\varepsilon) \log^2 n \cdot \log^2 N)$ update time for the unweighted and $O(\exp(1/\varepsilon) \log^2 n \cdot \log^5 N \cdot \log W)$ update time for the weighted case. They cast interval scheduling as the problem of finding a maximum independent set among a set of intervals on the x -axis. The authors extend this setting to multiple dimensions and design algorithms for approximating the maximum independent set among a set of

d -dimensional hypercubes, achieving a $(1 + \varepsilon)2^d$ -approximation in the unweighted and a $(4 + \varepsilon)2^d$ -approximation in the weighted regime.

The authors of [8] primarily focus on the unweighted case of approximating maximum independent set of a set of cubes. For the 1-dimensional case, which equals interval scheduling on one machine, they obtain $O(\log n/\varepsilon^2)$ update time, which is slower by a factor of $1/\varepsilon$ than our approach. They also show that their approach generalizes to the d -dimensional case, requiring poly log n amortized update time and providing $O(4^d)$ approximation.

The problem of dynamically maintaining an exact solution to interval scheduling on one or multiple machines is studied by [10]. They attain a guarantee of $\tilde{O}(n^{1/3})$ update time for unweighted interval scheduling on $M = 1$ machine, and $\tilde{O}(n^{1-1/M})$ for $M \geq 2$. Moreover, they show an almost-linear time conditional hardness lower bound for dynamically maintaining an exact solution to the weighted interval scheduling problem on even just $M = 1$ machine. This further motivates work such as ours that dynamically maintains approximate solutions for weighted interval scheduling.

The authors of [11] consider dynamic interval scheduling on multiple machines in the setting where all the jobs must be scheduled. The worst-case update time of their algorithm is $O(\log(n) + d)$, where d refers to the depth of what they call *idle intervals* (depth meaning the maximal number of intervals that contain a common point); they define an idle interval to be the period in a schedule between two consecutive jobs in a given machine. The same set of authors, in [12], also study dynamic algorithms for the monotone case, in which no interval completely contains another one. For this setup, they obtain an algorithm with $O(\log(n))$ update and query time.

In the standard model of computing (i.e., one processor, static), there exists an $O(n + m)$ running time algorithm for (exactly) solving the unweighted interval scheduling problem on a single machine with n jobs and integer coordinates bounded by m [13]. An algorithm with running time independent of m is described in [14], where it is shown how to solve this problem on M machines in $O(n \log(n))$ time. An algorithm is designed in [15] for weighted interval scheduling on M machines that runs in $O(n^2 \log(n))$ time.

We refer a reader to [3] and references therein for additional applications of the interval scheduling problem.

Other related work. There has also been a significant interest in job scheduling problems in which our goal is to schedule *all* the given jobs across multiple machines, with the objective to minimize the total scheduling time. Several variants have been studied, including setups that allow preemptions or setting where jobs have precedence constraints. We refer a reader to [16–22] and references therein for more details on these and additional variants of job scheduling. Beyond dynamic algorithms for approximating maximum independent sets of intervals or hypercubes, [23] show results for geometric objects such as disks, fat polygons, and higher-dimensional analogs. After we had published a preprint of this work, [23] proved a result for dynamic data structures approximating the maximum independent set of fat objects. As discussed in their Section 6, this subsumes our Theorem 1.1.

2 Problem Setup

In the interval scheduling problem, we are given n jobs and M machines. With each job j are associated two numbers s_j and $l_j > 0$, referring to “start” and “length” respectively, meaning that the job j takes l_j time to be processed and its processing can only start at time s_j . The job then finishes at time $f_j = s_j + l_j$. In addition, with each job j is associated weight/reward $w_j > 0$, that refers to the reward for processing the job j . The task of *interval scheduling* is to schedule jobs across machines while maximizing the total reward and respecting that each of the M machines can process at most one job at any point in time.

3 Overview of Our Techniques

Our primary goal is to present unified techniques for approximating scheduling problems that can be turned into efficient algorithms for many settings. In this section, we discuss key insights of our techniques.

In the problems our work tackles, partitioning the problem instance into independent, manageable chunks is crucial. Doing so enables an LCA to determine information about a job of interest without computing an entire schedule, or enables a dynamic data structure to maintain a solution without restarting from scratch.

3.1 Unweighted Interval Scheduling—Partitioning Over Time (Sect. 4)

For simplicity of presentation, we begin by examining our method for partitioning over time for just the unweighted interval scheduling problem on one machine (i.e., $M = 1$). In particular, we first focus on doing so for the dynamic setting.

Recall that in this setting, the primary motivation for partitioning over time is to divide the problem into independent, manageable chunks that can be utilized by a data structure to quickly modify a solution while processing an update. In our work, we partition the time dimension by maintaining a set of *borders* that divide time into some contiguous regions. By doing so, we divide the problem into many *independent regions*, and we ignore jobs that intersect multiple regions; equivalently, we ignore jobs that contain a border. Our goal is then to dynamically maintain borders in a way such that we can quickly recompute the optimal solution completely within some region, and that the suboptimality introduced by these borders does not affect our solution much. In Sect. 4, we show that by maintaining borders where the optimal solution inside each region, i.e., a time-range between two borders, is of size $\Theta(\frac{1}{\varepsilon})$, we can maintain a $(1 + \varepsilon)$ -approximation of an optimal solution as long as we optimally compute the solution within each region.

Here, the underlying intuition is that because each region has a solution of size $\Omega(\frac{1}{\varepsilon})$, we can charge any suboptimality caused by a border against the selected jobs in an adjacent region. Likewise, because each region’s solution has size $O(\frac{1}{\varepsilon})$, we are able to recompute the optimal solution within some region quickly using a balanced

binary search tree. We dynamically maintain borders satisfying our desired properties by adding a new border when a region becomes too large, or merging with an adjacent region when a region becomes too small. As only $O(1)$ regions will require any modification when processing an update, this method of partitioning time, while simple, enables us to improve the fastest known update/query time to $O(\log(n)/\varepsilon)$.¹ In Sect. 3.2 we build on these ideas to design an algorithm for the weighted interval scheduling problem.

3.2 Weighted Interval Scheduling (Sect. 5)

In our most technically involved result, we design the first deterministic $(1 + \varepsilon)$ approximation algorithm for weighted interval scheduling that runs in $\text{poly}(\log n, \frac{1}{\varepsilon})$ time. In this section, we give an outline of our techniques and discuss key insights. For full details, we refer a reader to Sect. 5.

3.2.1 Job Data Structure

Let \mathcal{E} be the set of all the endpoints of given jobs, i.e., \mathcal{E} contains s_i and f_i for each job $[s_i, f_i]$. We build a hierarchical data structure over \mathcal{E} as follows. This structure is organized as a binary search tree T . Each node Q of T contains a value $\text{KEY}(Q) \in \mathcal{E}$, with a “1-1” mapping between \mathcal{E} and the nodes of T . Each node Q is *responsible for a time range*. The root of T , that we denote by Q_{root} , is responsible for the entire time range $(-\infty, \infty)$. Each node Q has at most two children, that we denote by Q_L and Q_R . If Q is responsible for the time range $[X, Y]$, then Q_L is responsible for $[X, \text{KEY}(Q)]$, while Q_R is responsible for $[\text{KEY}(Q), Y]$.

Jobs are then assigned to nodes, where a job J is assigned to every node Q such that J is contained within the Q 's responsible time range.

3.2.2 Organizing Computation (Sect. 5.1)

We now outline how the structure T is used in computation. As a reminder, our main goal is to compute a $(1 + \varepsilon)$ -approximate weighted interval scheduling. This task is performed by requesting Q_{root} to solve the problem for the range $(-\infty, \infty)$. However, instead of computing the answer for the entire range $(-\infty, \infty)$ directly, Q_{root} *partitions* the range $(-\infty, \infty)$ into:

- at most $\text{poly}(n, 1/\varepsilon)$ ranges over which it is relatively easy to compute approximate solutions, such ones are called *sparse*, and
- at most $\text{poly}(n, 1/\varepsilon)$ remaining ranges over which it is relatively hard to compute approximate solutions at the level of Q_{root} .

These hard-to-approximate ranges are deferred to the children of Q_{root} , and are hard to approximate because any near-optimal solution for the range contains many jobs.

¹ The main advantage of this techniques is that it leads to worst-case $O(\log(n)/\varepsilon)$ update time, as opposed to only an amortized one. We point out that it is not difficult to obtain such amortized guarantee in the following way: after each $\varepsilon \cdot OPT$ many updates, recompute the optimum solution from scratch. Given access to the balanced binary tree structure described above, this re-computation can be done in $O(OPT \cdot \log n)$ time.

On the other hand, solutions in sparse ranges are of size $O(1/\varepsilon)$. As discussed later, approximate optimal solutions within sparse ranges can be computed efficiently; for details, see the paragraph *Approximate dynamic programming* below.

In general, a child Q_C of Q_{root} might receive *multiple* ranges from Q_{root} for which it is asked to find an approximately optimal solution. Q_C performs computation in the same manner as Q_{root} did – the cell Q_C partitions each range it receives into “easy” and “hard” to compute subranges. Q_C computes the first type of subranges, while the second type is deferred to the children of Q_C . These “hard” ranges have large weight and allow for drawing a boundary and hence dividing a range into two or more *independent* ranges. We now discuss how the partitioning into ranges is undertaken.

3.2.3 Auxiliary Data Structure (Sect. 5.2)

To divide a range into “easy” and “hard” ranges at the level of a node Q , we design an auxiliary data structure, which relates to a rough approximation of the problem. This structure, called $Z(Q)$, maintains a set of points (we call these points *grid endpoints*) that partition Q into *slices of time*. We use slice to refer to a time range between two *consecutive* points of $Z(Q)$. Recall how for unweighted interval scheduling, we maintained a set of borders and ignored a job that crossed any border. In the weighted version, we will instead use $Z(Q)$ as a set of partitions from which we will use *some subset* to divide time. Our method of designing $Z(Q)$ reduces the task of finding a partitioning over time $Z(Q)$ within a cell for the $(1 + \varepsilon)$ -approximate weighted interval scheduling problem to finding multiple partitionings for the $(1 + \varepsilon)$ -approximate unweighted problem.

It is instructive to think of $Z(Q)$ in the following way. First, we view weighted interval scheduling as $O(\log n)$ independent instances of unweighted interval scheduling – instance i contains the jobs having weights in the interval $(w_{max}(Q)/2^{i+1}, w_{max}(Q)/2^i]$. Then, for each unweighted instance we compute borders as described in Sect. 3.1. $Z(Q)$ constitutes a subset of the union of those borders across all unweighted instances. We point out that the actual definition of $Z(Q)$ contains some additional points that are needed for technical reasons, but in this section we will adopt this simplified view. In particular, as we will see, $Z(Q)$ is designed such that the optimal solution within each slice has small total reward compared to the optimal solution over the entirety of Q . This enables us to partition the main problem into subproblems such that the suboptimality of discretizing the time towards slices, that we call *snapping*, is negligible.

However, a priori, it is not even clear that such structure $Z(Q)$ exists. So, one of the primary goals in our analysis is to show that there exists a near-optimal solution of a desirable structure that can be captured by $Z(Q)$. The main challenge here is to detect/localize sparse and dense ranges efficiently and in a way that yields a fast dynamic algorithm. As an oversimplification, we define a solution as having *nearly-optimal sparse structure* if it can be generated with roughly the following process:

- Each cell Q receives a set of disjoint time ranges for which it is supposed to compute an approximately optimal solution using jobs assigned to Q or its descendants. Each received time range must have starting and ending time in $Z(Q)$.




Fig. 1 Visual example for hierarchical decomposition. Consider we are given jobs with the following ranges of $(1, 5), (2, 10), (7, 20), (4, 5)$. On the left is T , a balanced binary search tree over the set of all s_i and f_i . On the right is the hierarchical decomposition that corresponds to T . That is, in each row, the intervals on the right correspond to the $[l_Q, r_Q]$ for the nodes on the left. For instance, in the third row, $(-\infty, 2]$ corresponds to the node Q with $KEY(Q) = 1$

- For each time range \mathcal{R} that Q receives, the algorithm partitions \mathcal{R} into disjoint time ranges of three types: sparse time ranges, time ranges to be sent to Q_L for processing, and time ranges to be sent to Q_R for processing. In particular, this means that subranges of \mathcal{R} are deferred to the children of Q for processing.
- For every sparse time range, Q computes an optimal solution using at most $1/\epsilon$ jobs.
- The union of the reward/solution of all sparse time ranges on all levels must be a $(1 + \epsilon)$ -approximation of the globally optimal solution without any structural requirements.

Moreover, we develop a *charging method* that enables us to partition each cell with only $|Z(Q)| = \text{poly}(1/\epsilon, \log(n))$ points and still have the property that it contains a $(1 + \epsilon)$ -approximately optimal solution with nearly-optimal sparse structure. Then, we design an approximate dynamic programming approach to efficiently compute near-optimal solutions for sparse ranges. Combined, this enables a very efficient algorithm for weighted interval scheduling. On a high-level, $Z(Q)$ enables us to eventually decompose an entire solution into sparse regions.

3.2.4 The Charging Method (Sect. 5.2.3)

We now outline insights of our charging arguments that enable us to convert an optimal solution OPT into a near-optimal solution OPT' with nearly-optimal sparse structure while relaxing our partitioning to only need $|Z(Q)| = \text{poly}(1/\epsilon, \log(N))$ points. For a visual aid, see Fig. 2.

As outlined in our overview of the nearly-optimal sparse structure, each cell Q receives a set of disjoint time ranges, with each time range having endpoints in $Z(Q)$, and must split them into three sets: sparse time ranges, time ranges for Q_L , and time

ranges for Q_R . We will now modify OPT by deleting some jobs. This new solution will be denoted by OPT' and will have the following properties:

1. OPT' exhibits nearly-optimal sparse structure; and
2. OPT' is obtained from OPT by deleting jobs of total reward at most $O(\varepsilon \cdot w(OPT))$.

We outline an example of one such time range a cell Q may receive in Fig. 2, annotated by “received range \mathcal{R} ”. We will color jobs in Fig. 2 to illustrate aspects of our charging argument, but note that jobs do not actually have a color property beyond this illustration. Since our structure only allows a cell Q to use a job within its corresponding time range, any relatively valuable job that crosses between Q_L and Q_R must be used now by Q putting it in a sparse time range. One such valuable job in Fig. 2 is in blue marked by “B”. To have “B” belong to a sparse range, we must divide the time range \mathcal{R} somewhere, as otherwise our solution in the received range will be dense. If we naively divide \mathcal{R} at the partition of $Z(Q)$ to the left and right of the job “B”, we might be forced to delete some valuable jobs; such jobs are pictured in green and marked by “G”. Instead, we expand the division outwards in a more nuanced manner. Namely, we keep expanding outwards and looking at the job that contains the next partition point (if any). If the job’s value exceeds a certain threshold, as those pictured as green and marked by “G” in Fig. 2, we continue expanding. Otherwise, the job crossing a partition point is below a certain threshold, pictured as brown and not marked in Fig. 2, and its deletion can be charged against the blue job. We delete such brown jobs, and the corresponding partition points, i.e., the vertical red lines crossing those brown jobs, constitute the start and the end of the sparse range. By the end, we decided the starting and ending time of the sparse range, and what remains inside are blue job(s), green job(s), and yellow job(s) (also marked by “Y”). Note that yellow jobs must be completely within a partition slice of $Z(Q)$. Since we define $Z(Q)$ such that the optimal total reward within any grid slice is small, the yellow jobs have relatively small rewards compared to the total reward of green and blue jobs that we know must be large. Accordingly, we can delete the yellow jobs (to help make this time range’s solution sparse) and charge their cost against a nearby green or blue job. In Fig. 2, an arrow from one job to another represents a deleted job pointing towards the job which we charge its loss against. Finally, each sparse range contains only green job(s) and blue job(s). If there are more than $1/\varepsilon$ jobs in such a sparse range, we employ a simple sparsifying step detailed in the full proof.

It remains to handle the time ranges of the received range that were not put in sparse ranges. These time ranges are sent to Q_L and Q_R . In Fig. 2, these ranges are outlined in yellow and annotated by “child subproblem”. However, the time ranges do not necessarily align with $Z(Q_L)$ or $Z(Q_R)$ as is required by nearly-optimal sparse structure. We need to adjust these ranges to align with $Z(Q_L)$ or $Z(Q_R)$ so we can send the ranges to the children. See Fig. 3 for intuition on why we cannot just immediately “snap” these child subproblems to the partition points in $Z(Q_L)$ and $Z(Q_R)$. (We say that a range \mathcal{R} is *snapped* inward (outward) within cell Q if \mathcal{R} is shrunk (extended) on both sides to the closest points in $Z(Q)$). Inward snapping is illustrated in Fig. 3.) Instead, we employ a similar charging argument to deal with snapping. As an analog to how we expanded outwards from the blue job for defining sparse ranges,




Fig. 2 Visual example for charging argument

we employ a charging argument where we contract inwards from the endpoints of the child subproblem. In summary, these charging arguments enabled us to show a solution of nearly-optimal sparse structure exists even when only partitioning each cell Q with $|Z(Q)| = \text{poly}(1/\epsilon, \log(n))$ points.

3.2.5 Approximate Dynamic Programming (Sect. 5.3)

Now, we outline our key advance for more efficiently calculating the solution of nearly-optimal sparse structure. This structure allows us to partition time into ranges with sparse solutions. More formally, we are given a time range and we want to approximate an optimal solution within that range that uses at most $1/\epsilon$ jobs. We outline an approximate dynamic programming approach that only requires polynomial time dependence on $1/\epsilon$.

The relatively well-known dynamic programming approach for computing weighted interval scheduling is maintaining a dynamic program where the state is a prefix range of time, and the output is the maximum total reward that can be obtained in that range of time. However, for our purposes, there are too many possibilities for prefix ranges of time to consider. Instead, we invert the dynamic programming approach and have a state referencing some amount of reward, where the dynamic program returns the minimum length prefix range of time in which one can obtain a given reward. Unfortunately, there are also too many possible amounts of rewards. We observe that we do not actually need this exact state but only an approximation. In particular, we show that one can round this state down to powers of $(1 + \epsilon^2)$ and hence significantly reduce the state space. In Sect. 5.3, we show how one can use this type of observation to quickly compute approximate dynamic programming for a near-optimal sparse solution inside any time range.

3.2.6 Comparison with Prior Work

The closest to our work is the one of [9]. In terms of improvements, we achieve the following: we remove the dependence on N and w_{\max} in the running-time analysis; and, we design an algorithm with $\text{poly}(1/\epsilon, \log n)$ update/query time, which is exponentially faster in $1/\epsilon$ compared the prior work.

In this prior work, jobs are assumed to have lengths of at least 1 and belong in the time-interval $[1, N]$. To remove the dependence on N and such assumptions, we designed a new way of bookkeeping jobs. Instead of using a complete binary tree on $[1, N]$ to organize jobs as done in the prior work, we construct a binary balanced search tree on the endpoints of jobs. A complete binary tree on $[1, N]$ is oblivious to the density of jobs. On the other hand, and intuitively, our approach allows for “instance-based” bookkeeping: the jobs are in a natural way organized with respect to their density. Resorting to this approach incurs significant technical challenges. Namely, the structure of the solution our tree maintains is hierarchically organized. However, each tree update, which requires node-rotations, breaks this structure which requires additional care in efficiently maintaining approximate solution after an update, as well as requiring an entirely different approach for maintaining a partitioning of time $Z(Q)$ within cells. Moreover, we show how to leverage these ideas further to obtain a deterministic approach.

In our work, we use borders to define the so-called sparse and dense ranges. This idea is inspired by the work of [9]. We emphasize, though, that one of our main contributions and arguably the most technically involved component is showing how to algorithmically employ those borders in running-time only polynomially dependent on $1/\varepsilon$, while [9] require exponential dependence on $1/\varepsilon$.

Our construction of auxiliary data structure $Z(Q)$ enables us to boost an $O(\log(n))$ -approximate solution into a decomposition enabling a $(1 + \varepsilon)$ -approximate solution is inspired by the approach of [9]. They similarly develop $Z(Q)$ to boost an instead $O(1)$ -approximation that fundamentally relies on the bounded coordinate assumptions of jobs being within $[1, N]$ and having length at least 1. Our different approach towards $Z(Q)$ enables simplification of some arguments, or on length or bounded coordinate assumptions. Further, we note that the dynamic programming approach for sparse regions that we develop is significantly faster than the enumerative approach used in the prior work, that eventually enables us to obtain a $\text{poly}(1/\varepsilon)$ dependence in the running time. The way we combine solutions over sparse regions is similar to the way it is done in the prior work.

4 Dynamic Unweighted Interval Scheduling on a Single Machine

In this section, we prove Theorem 1.1. As a reminder, Theorem 1.1 considers the case of interval scheduling in which $w_j = 1$ for each j and $M = 1$, i.e., the jobs have unit reward and there is only a single machine at our disposal. This case can also be seen as a task of finding a maximum independent set among intervals lying on the x -axis. The crux of our approach is in designing an algorithm that maintains the following invariant:

Invariant 1 *The algorithm maintains a set of borders such that an optimal solution schedules between $1/\varepsilon$ and $2/\varepsilon$ intervals within each two consecutive borders.*

We will maintain this invariant unless the optimal solution has fewer than $1/\varepsilon$ intervals, in which case we are able to compute the solution from scratch in negligible time. We aim for our algorithm to maintain Invariant 1 while keeping track of the optimal

solution between each pair of consecutive borders. The high-level intuition for this is that if we do not maintain too many borders, then our solution must be very good (our solution decreases in size by at most 1 every time we add a new border). Furthermore, if the optimal solution within borders is small, it is likely easier for us to maintain said solutions. We prove that this invariant enables a high-quality approximation:

Lemma 4.1 *A solution that maintains an optimal solution within consecutive pairs of a set of borders, where the optimal solution within each pair of consecutive borders contains at least K intervals, maintains a $\frac{K+1}{K}$ -approximation.*

Proof For our analysis, suppose there are implicit borders at $-\infty$ and $+\infty$ so that all jobs are within the range of borders. Consider an optimal solution OPT . We will now design a K -approximate optimal solution OPT' as follows: given OPT , delete all intervals in OPT that overlap a drawn border. Fix an interval J appearing in OPT but not in OPT' . Assume that J intersects the i -th border. Recall that between the $(i - 1)$ -st and the i -th border there are at least K intervals in OPT' . Moreover, at most one interval from OPT intersects the i -th border. Hence, to show that OPT' is a $\frac{K+1}{K}$ -approximation of OPT , we can charge the removal of J to the intervals appearing between the $(i - 1)$ -st and the i -th border in OPT' . \square

Not only does Invariant 1 enable high-quality solutions, but it also assists us in quickly maintaining such a solution. We can maintain a data structure with $O(\frac{\log(n)}{\varepsilon})$ updates and $O(\log(n))$ queries that moves the borders to maintain the invariant and thus maintains an $(1 + \varepsilon)$ -approximation as implied by Lemma 4.1.

Theorem 1.1 (Unweighted dynamic, single machine) *Let \mathcal{J} be a set of n jobs. For any $\varepsilon > 0$, there exists a fully dynamic algorithm for $(1 + \varepsilon)$ -approximate unweighted interval scheduling for \mathcal{J} on a single machine performing updates in $O\left(\frac{\log(n)}{\varepsilon}\right)$ and queries in $O(\log(n))$ worst-case time.*

Proof Our goal now is to design an algorithm that maintains Invariant 1, which by Lemma 4.1 and for $K = 1/\varepsilon$ will result in a $(1 + \varepsilon)$ -approximation of MAXIMUM-IS.

On a high level, our algorithm will maintain a set of borders. When compiling a solution of intervals, the algorithm will not use any interval that contains any of the borders but proceed by computing an optimal solution between each two consecutive borders. The union of those between-border solutions is the final solution. Moreover, we will maintain the invariant that the optimal solution for every contiguous region is of size within $[\frac{1}{\varepsilon}, \frac{2}{\varepsilon})$.

In the rest, we show how to implement these steps in the claimed running time. *Maintained data-structures.* Our algorithm maintains a balanced binary search tree T_{all} of intervals sorted by their starting points. Each node of T_{all} will also maintain the end-point of the corresponding interval. It is well-known how to implement a balanced binary search tree with $O(\log n)$ worst-case running time per insertion, deletion, and search query. Using such an implementation, the algorithm can in $O(\log n)$ time find the smallest ending point in a prefix/suffix on the intervals sorted by their starting points. That is, in $O(\log n)$ time we can find the interval that ends earliest, among those that start after a certain time.

In addition, the algorithm also maintains a balanced binary search tree T_{borders} of the borders currently drawn.

Also, we will maintain one more balanced binary search tree T_{sol} that will store the intervals that are in our current solution.

We will use that for any range with an optimal solution of size S , we can make $O(S)$ queries to these data structures to obtain an optimal solution for the range in $O(S \cdot \log n)$ time.

Update after an insertion. Upon insertion of an interval J , we add J to T_{all} . We make a query to T_{borders} to check whether J overlaps a border. If it does, we need to do nothing; in this case, we ignore J even if it belongs to an optimal solution. If it does not, we recompute the optimal solution within the two borders adjacent to J . If after recomputing, the new solution between the two borders is too large, i.e., it has at least $\frac{2}{\varepsilon}$ intervals, then draw/add a border between the $\frac{1}{\varepsilon}$ -th and the $(1 + \frac{1}{\varepsilon})$ -th of those intervals.

Update after a deletion. Upon deletion of an interval J , we delete J from T_{all} . If J was not in our solution, we do nothing else. Otherwise, we recompute the optimal solution within the borders adjacent to J and modify T_{sol} accordingly. Let those borders be the i -th and the $(i + 1)$ -st. If the new solution between borders i and $i + 1$ now has a size less than $1/\varepsilon$ (it would be size exactly $1/\varepsilon$), we delete an arbitrary one of the two borders (thus combining this region with an adjacent region). Then, we recompute the optimal solution within the (now larger) region J is in. If this results in a solution of size at least $2/\varepsilon$, we will need to split the newly created region by adding a border. Before splitting, the solution will have size upper-bounded by one more than the size of the solutions within the two regions before combining them as an interval may have overlapped the now deleted border (one region with size exactly $\frac{1}{\varepsilon} - 1$ and the other upper-bounded by $\frac{2}{\varepsilon} - 1$). Thus, the solution has size at in range $[2/\varepsilon, \frac{3}{\varepsilon})$. We can add a border between interval $1/\varepsilon$ and $1/\varepsilon + 1$ of the optimal solution and will have a region with exactly $1/\varepsilon$ intervals and another with $[1/\varepsilon, 2/\varepsilon)$ intervals, maintaining our invariant.

In all of these, the optimal solution for each region has size $O(1/\varepsilon)$, so recomputing takes $O(\log(n)/\varepsilon)$ time.

For queries, we will have maintained T_{sol} in our updates such that it contains exactly the intervals in our solution. So each query we just need to do a lookup to see if the interval is in T_{sol} in $O(\log n)$ time. \square

This result improves the best-known time complexities [8, 9]. Unfortunately, it does not immediately generalize well to the weighted variant. In Sect. 5, we show our more technically-challenging result for the weighted variant.

5 Dynamic Weighted Interval Scheduling on a Single Machine

This section focuses on a more challenging setting in which jobs have non-uniform weights. Non-uniform weights introduce difficulties for the approach mentioned in Sect. 4, as adding a border (which entails ignoring all the jobs that cross that border) may now force us to ignore a very valuable job. Straightforward extensions of this

border-based approach require at least a linear dependence on the ratio between job rewards (e.g., if all jobs have rewards within $[1, w]$, then straightforward extensions would require a linear dependence on w). This is because an ignored job containing a border can have a reward of w (as opposed to just 1), requiring w/ε reward inside the region to charge it against (as opposed to just $1/\varepsilon$). In this work, we show how to perform this task in $O(\text{poly}(\log(n), 1/\varepsilon))$ time, having no such dependency on the rewards of the jobs or the starting/ending times. This improves upon the best-known preexisting result of $O(\text{poly}(\log(n), \log(N), \log(w)) \cdot \exp(1/\varepsilon))$ time accomplished by the decomposition scheme designed in the work of Henzinger et al. [9], which we compare with in Sect. 3.2.6. Both our algorithm and our analysis introduce new ideas that enable us to design a dynamic algorithm with running time having only polynomial dependence on $1/\varepsilon$ and $\log(n)$, yielding an exponential improvement in terms of $1/\varepsilon$ over [9], and removing all dependence on N and w . Moreover, our algorithm is deterministic and requires no assumption on the lengths or coordinate values of the jobs; [9] is also deterministic, but it assumes all jobs are length at least 1 and all coordinates are within $[0, N]$, where N affects the time complexity.

As the first step, we show that there exists a $(1 + \varepsilon)$ -approximate optimal solution OPT' that has *nearly-optimal sparse structure*, similar to a structure used in [9]. We define properties of this structure in Sect. 5.2, although it is instructive to think of this structure as of a set of non-overlapping time ranges such that:

1. Within each time range, there is an approximately optimal solution that contains a small number of jobs (called *sparse*);
2. The union of solutions across all the time ranges is $(1 + \varepsilon)$ -approximate; and
3. There is an efficient algorithm to obtain these time ranges.

Effectively, this structure partitions time such that we get an approximately optimal solution by computing sparse solutions within partitioned time ranges and ignoring jobs that are not fully contained within one partitioned time range. To obtain the guarantees of such a set of time ranges that can be obtained efficiently, we utilize a new hierarchical decomposition based on a balanced binary search tree and employ novel charging arguments. This result is described in detail in Sect. 5.2.

Once equipped with this structural result, we first design a dynamic programming approach to compute an approximately optimal solution within one time range. Let w_{max} denote the maximal reward among all jobs currently in the instance. To obtain an algorithm whose running time is proportional to the number of jobs in the solution for a time range, as opposed to the length of that range, we “approximate” states that our dynamic programming approach maintains, and ultimately obtain the following claim whose proof is deferred to Sect. 5.3.

Lemma 5.1 *Given any contiguous time range \mathcal{R} and an integer K , consider an optimal solution $OPT(\mathcal{R}, K)$ in \mathcal{R} containing at most K jobs and ignoring jobs with weight less than $\varepsilon/n \cdot w_{max}$. Then, there is an algorithm that in \mathcal{R} finds a $(1 + \varepsilon)$ -approximate solution to $OPT(\mathcal{R}, K)$ in $O\left(\frac{K \log(n) \log^2(K/\varepsilon)}{\varepsilon^2}\right)$ time and with at most $O\left(\frac{K \log(K/\varepsilon)}{\varepsilon}\right)$ jobs.*

Observe that running time of the algorithm given by Lemma 5.1 has no dependence on the length of \mathcal{R} . Also observe that the algorithm possibly selects slightly more than

K jobs to obtain a $(1 + \varepsilon)$ -approximation of the best possible reward one could obtain by using at most K jobs in \mathcal{R} (i.e., $OPT(\mathcal{R}, K)$).

Finally, in Sect. 5.5 we combine all these ingredients and prove the main theorem of this section.

Theorem 1.2 (Weighted dynamic, single machine) *Let \mathcal{J} be a set of n weighted jobs. For any $\varepsilon > 0$, there exists a fully dynamic algorithm for $(1 + \varepsilon)$ -approximate weighted interval scheduling for \mathcal{J} on a single machine performing updates and queries in worst-case time $T \in \text{poly}(\log n, \frac{1}{\varepsilon})$. The exact complexity of T is given by*

$$O\left(\frac{\log^{12}(n)}{\varepsilon^7} + \frac{\log^{13}(n)}{\varepsilon^6}\right).$$

5.1 Decomposition Overview

We utilize a hierarchical decomposition to organize time such that we may efficiently obtain time ranges that satisfy the nearly-optimal sparse structure. This decomposition has two levels of granularity. For the higher-level decomposition, we employ a decomposition similar to that of a balanced binary search tree with $O(\log(n))$ depth. Each cell Q in this balanced binary search tree will correspond to a range of time. Further details on this hierarchical decomposition are described in Sect. 5.2.1.

For the lower-level decomposition, we split each cell Q more finely. Formally, for a set of grid endpoints $Z(Q)$, we define a grid slice as follows.

Definition 5.2 (*Grid slice*) Given a set of grid endpoints $Z(Q) = \{r_1, r_2, \dots, r_{X-1}\}$ with $r_i < r_{i+1}$, we use *grid slice* to refer to an interval (r_i, r_{i+1}) , for any $1 \leq i < X - 1$. Note that a grid slice between r_i and r_{i+1} does not contain r_i nor r_{i+1} .

We further discuss $Z(Q)$ in Sect. 5.2.2. Importantly, $Z(Q)$ is designed such that the optimal solution entirely within any grid slice is upper-bounded to be relatively small compared to the weight of the optimal solution within Q , or $w(OPT(Q))$. This property makes the grid endpoints $Z(Q)$ a helpful tool in partitioning time. At a high level, $Z(Q)$ is used to define a set of segments that motivate dynamic programming states of the form $DP(Q, S)$, where each S corresponds to a segment between two grid endpoints of $Z(Q)$, and $DP(Q, S)$ computes an approximately optimal sparse solution among schedules that can only use jobs contained within the segment of time S . The key idea is that this dynamic programming enables the partitioning of time into *dense* and *sparse* ranges. Solutions for sparse ranges are computed immediately, while dense ranges are solved by children with dynamic programming (by further dividing the dense range into more sparse and dense ranges). We recall from Sect. 3.2.6 that [9] were first to design a two-level hierarchical decomposition that computes $DP(Q, S)$ to optimize over dense and sparse ranges. However, we emphasize that our work utilizes entirely new approaches for our high-level hierarchical decomposition into cells Q , for our low-level decomposition of each cell into $Z(Q)$, and for our method of computing approximately optimal sparse solutions of $DP(Q, S)$.

5.2 Solution of Nearly-Optimal Sparse Structure

To remove exponential dependence on $1/\varepsilon$ and all dependence on N and w , we introduce a new algorithm for approximating sparse solutions, a new hierarchical decomposition, and novel charging arguments that (among other things) reduce the number of grid endpoints $|Z(Q)|$ required in each cell. With this, we will compute an approximately optimal solution of the following very specific structure.

Definition 5.3 (*Nearly-optimal sparse structure*) To have nearly-optimal sparse structure, a solution must be able to be generated with the following specific procedure:

- Each cell Q will receive a set of time ranges, denoted as $RANGES(Q)$, with endpoints in $Z(Q)$. To start, Q_{root} will receive one time range containing all of time (i.e., $RANGES(Q_{root}) = \{[-\infty, \infty]\}$)
- $RANGES(Q)$ is split into a collection of disjoint time ranges, with each being assigned to one of three sets: $SPARSE(Q)$, $RANGES(Q_L)$, $RANGES(Q_R)$
- $SPARSE(Q)$, a set of time ranges, must have endpoints in $Z(Q) \cup Z(Q_L) \cup Z(Q_R)$
- For each child Q_{child} (where $child \in \{L, R\}$) of Q , $RANGES(Q_{child})$ must have all endpoints in $Z(Q_{child})$
- The total weight of sparse solutions (solutions with at most $1/\varepsilon$ jobs) within sparse time ranges must be large (where $SPARSE_OPT(\mathcal{R})$ denotes an optimal solution having at most $1/\varepsilon$ jobs within range \mathcal{R}):

$$\sum_Q \sum_{R \in SPARSE(Q)} w(SPARSE_OPT(\mathcal{R})) \geq (1 - O(\varepsilon))w(OPT)$$

Now, we prove our result for a $(1 + \varepsilon)$ -approximation to dynamic, weighted interval MAXIMUM- IS algorithm with only polynomial time dependence on $1/\varepsilon$ and $\log(n)$. Unlike the decomposition of Henzinger et al., we will not define our decomposition such that each cell Q will split exactly in half to produce both its children Q_L and Q_R . Instead, we will divide every cell Q in a manner informed by a balanced binary search tree. Desirably, this will make the depth of our decomposition $O(\log(n))$ instead of $O(\log(N))$, but it will remove the possibility of utilizing the random-offset style of idea to assign jobs to cells where they each job's length is approximately ε fraction of the cell's length. This necessitates novel charging arguments. We supplement this new hierarchical decomposition with a new alternative for the $Z(Q)$ data structure that enables us to determine important dynamic program subproblems without any dependence on N . Additionally, we take a new approach for solving the small sparse subproblems, where we use an approximate dynamic programming idea to remove exponential dependence on $1/\varepsilon$ in the best known running time for these subproblems. In our novel charging arguments, there is a particular focus on changing where deleted intervals' weights are charged against and introducing a *snapping budget*, which we use to relax the required number of grid endpoints $|Z(Q)|$ to depend only polynomially on $1/\varepsilon$. As a reminder, $Z(Q)$ is a set of grid points within Q such that between any two consecutive points we are guaranteed that the optimal solution has small weight. Our final algorithm will consider a number of subproblems for each cell proportional to

$|Z(Q)|^2$, so improvements in $|Z(Q)|$ directly lead to improvements in the best-known running time. Effectively, we make each of our smaller subproblems easier to solve while also reducing the number of subproblems we need to solve. All improvements are exponential in ε and remove dependence on N and w .

5.2.1 Hierarchical Decomposition

We now formally describe our hierarchical decomposition of jobs.

- Consider the set of all jobs’ starting/ending times, i.e., for each job i , include s_i and f_i . Now, consider a balanced binary search tree T over this set of times. For the sake of this paper, one can assume this is maintained by a red-black tree such that the tree has depth $O(\log(n))$ and $O(\log(n))$ rotations are required per update. We have a cell Q in our hierarchical decomposition corresponding to each node in T . Let $KEY(Q)$ be the corresponding key for the node in T .
- Each Q has a left child Q_L or right child Q_R if the corresponding node in T does.
- Each cell Q represents a range of time. Q_{root} corresponds to all time, meaning $TIME(Q_{root}) = [-\infty, \infty]$. This time range is split for the children of Q by $KEY(Q)$. More formally, given a cell Q where $TIME(Q) = [l_Q, r_Q]$, then (if Q_L exists) $TIME(Q_L) = [l_Q, KEY(Q)]$, and (if Q_R exists) $TIME(Q_R) = [KEY(Q), r_Q]$.

This fully describes our hierarchical decomposition of depth $O(\log(n))$. A visual example is provided in Fig. 1.

5.2.2 Structure $Z(Q)$

We use the set of grid points $Z(Q)$ to determine segments that will be used as subproblems for dynamic programming and in reference to the nearly-optimal sparse structure. For some specified X , our goal is to maintain a $Z(Q)$ such that the optimal solution within every grid slice is at most $O(w(OPT(Q))/X)$. The previously-utilized methods for obtaining this require logarithmic dependence on N and w . To remove dependence on w , we relax our requirements of $Z(Q)$ to ignore all jobs with weight less than $w(OPT(Q)) \cdot \varepsilon/n$; in total, these jobs have negligible reward. To remove dependence on N , we consider an alternative approach to computing $Z(Q)$, where we take the union of multiple solutions to $Z(Q)$ for the analogous unweighted interval scheduling problem using ideas similar to those in Sect. 4. We design a $Z(Q)$ with the following guarantees, whose proof is deferred to Sect. 5.4:

Lemma 5.4 (Dynamically maintaining $Z(Q)$) *For any fixed positive integer X , it is possible to return a set $Z(Q)$ for any cell Q in the hierarchical decomposition in $O(X \cdot \log^3(n))$ query time. Moreover, the returned $Z(Q)$ will satisfy the following properties:*

- For every Q , the optimal solution within each grid slice of $Z(Q)$ is at most $O(w(OPT(Q))/X)$; as a reminder, we ignore jobs with weights less than $w(OPT(Q)) \cdot \varepsilon/n$.
- For every Q , $|Z(Q)| = O(X \cdot \log^2(n))$




Fig. 3 This example illustrates why the snapping we perform must be done carefully. The horizontal segments in this figure represent jobs. We show an initial dense range (outlined in purple) with endpoints in $Z(Q)$. We show where these endpoints are in Q_L with dashed vertical lines. Importantly, they are not aligned with $Z(Q_L)$, i.e., the vertical dashed lines do not belong to $Z(Q_L)$. However, our structure *requires* that dense ranges align with $Z(Q_{child})$, so we must address this. If we were to naively snap the endpoints of the dense range inwards to the endpoints of $Z(Q_L)$, then we would need to delete some jobs (these deleted jobs are colored in yellow and marked by “Y”), while some other jobs would not be affected (like the remaining jobs in this example, those colored in blue). While this naive snapping may be fine in some cases, it will incur a significant loss in cases where the “Y” jobs have large weights. Notice that naively snapping outward to define a new region corresponding to the purple one is not a solution either, as this could cause the dense time range to overlap with a previously selected sparse time range. Having overlapping ranges can cause us to choose intersecting jobs and, thus, an invalid solution. Thus, we detail a more comprehensive manner of dealing with snapping

5.2.3 Existence of Desired $(1 + \epsilon)$ -Approximate Solution

We now argue that there exists a $(1 + O(\epsilon))$ -approximation with nearly-optimal sparse structure in reference to our new hierarchical decomposition for Q and our $Z(Q)$ when using $X = \frac{\log^2(n)}{\epsilon^2}$ and thus $|Z(Q)| = O(\frac{\log^4(n)}{\epsilon^2})$:

Lemma 5.5 *There exists a solution OPT' that has nearly-optimal sparse structure and such that $w(OPT') \geq (1 - O(\epsilon))w(OPT)$. Thus, OPT' is a $(1 + O(\epsilon))$ -approximation of OPT .*

Proof We emphasize that the goal of this lemma is not to show how to construct a solution algorithmically, but rather to show that there exists one, that we refer to by OPT' , that has a specific structure and whose weight is close to OPT .

In this paragraph, we provide a proof overview. At a high-level, we show this claim by starting with OPT , and maintaining a solution OPT' that holds our desired structure and only deletes jobs with total weight $O(\epsilon \cdot w(OPT))$. Our process of converting OPT to OPT' is recursive, as we start at the root and work down. Generally, our preference for any range $\mathcal{R} \in RANGES(Q)$ will be to defer it to a child by passing it on to a $RANGE(Q_{child})$. This preference can often not be immediately satisfied for two reasons: (i) \mathcal{R} may not be completely contained within a Q_{child} (i.e. \mathcal{R} crosses

between Q_L and Q_R), or (ii) the endpoints of \mathcal{R} do not align with the corresponding $Z(Q_{child})$. We will modify OPT to accommodate these concerns. To handle concern (i), we will delete a job in OPT if it crosses between Q_L and Q_R and has small weight (and hence it can be ignored). Otherwise, if such a crossing job has large weight, we will divide \mathcal{R} into three time ranges such that one is contained within Q_L , one uses the crossing job, and the last is contained within Q_R , using a process detailed in the following proof. For the central third, we will sparsify this range to produce a set $SPARSE(Q)$ of sparse time ranges. For time ranges completely contained within Q_L and Q_R that are not designated as sparse time ranges, we will essentially consider them dense time ranges, that will be delegated to children cells of Q . In order to delegate a time range to a child Q_{child} , we require that the delegated time range must have endpoints that align with $Z(Q_{child})$. Accordingly, we perform modifications to “snap” the time ranges’ endpoints to $Z(Q_{child})$ for the corresponding child Q_{child} of Q and include the “snapped” time ranges in $RANGES(Q_{child})$. We show that throughout this process, we do not delete much weight from OPT and obtain an OPT' that has our desired structure. Now, we present the proof in detail:

Deleting light crossing jobs.

We now describe how to modify OPT , obtaining OPT' , such that OPT' has our desired structure and OPT' is a $(1 + \epsilon)$ -approximation of OPT . Note that we will never actually compute OPT' . It is only a hypothetical solution that has nice structural properties and that we use to compare our output to.

For a cell Q , consider a time range it receives in $RANGES(Q)$. We shall split this time range into sparse time ranges (to be added to $SPARSE(Q)$) and dense time ranges (to be added to $RANGES(Q_L)$ or $RANGES(Q_R)$). There is at most one range $\mathcal{R}_{cross} \in RANGES(Q)$ that crosses between Q_L and Q_R , and we call the at most one job crossing between Q_L and Q_R the *crossing job* (if it exists). If the crossing job has weight $\leq \frac{\epsilon}{\log(n)} w(OPT(Q))$, we call it *light*, we delete the light crossing job, and we split \mathcal{R}_{cross} at the dividing point $KEY(Q)$. One of these two resulting ranges can inherit the snapping budget of \mathcal{R}_{cross} , while we can allocate the other a snapping budget of weight $O(\frac{\epsilon}{\log(n)} w(OPT(Q)))$. We delete/allocate at most $O(\frac{\epsilon}{\log(n)} w(OPT(Q)))$ weight at every cell, $O(\frac{\epsilon}{\log(n)} w(OPT))$ weight at every level, and $O(\epsilon w(OPT))$ weight in total. Also note how all ranges in $RANGES(Q)$ are now completely contained within either Q_L or Q_R . Otherwise, if the crossing job has large weight, we call it *heavy* and must find some way to include it in our solution instead of deleting it.

Utilizing heavy crossing jobs.

We now focus on showing how to construct our solution using a *heavy* crossing job. Our goal is to split \mathcal{R}_{cross} into three parts: one range completely within Q_L , some sparse ranges that will be $SPARSE(Q)$ and include the crossing job among other jobs, and one range completely within Q_R . As an overview, we will start by considering the smallest time range that contains the crossing job and spans the grid between two (not necessarily consecutive) endpoints in $Z(Q)$. This range may contain many jobs in OPT , so we perform an additional refinement to divide it up into sparse time ranges. In this refinement, we will split up the time range such that we do not delete too much weight and, moreover, all of the resulting time ranges have at most $1/\epsilon$ jobs. These time




Fig. 4 This example illustrates how the sparse regions are created. All vertical segments within Q , which are red in the figure, correspond to the points in $Z(Q)$. The cell Q is divided by $Z(Q)$ such that the optimal solution within every grid slice is small. As a reminder, a grid slice is an open time-interval between two consecutive points in $Z(Q)$; see Definition 5.2 for a formal definition. We start with the heavy crossing job (the blue horizontal segment marked by “B”). From this heavy crossing job, we expand the region outwards as necessary. In this example, we expanded to the right, seeing two highlighted jobs (the green horizontal segments marked by “G”) until we saw a job with low enough weight intersecting a grid endpoint (these job segments are colored in brown and crossed). We delete such brown jobs, and use the grid endpoints they intersected to define the region (outlined in purple and annotated by “new region”). Useless jobs (pictured in yellow) are then deleted. Later, we sparsify the region

ranges now constitute $SPARSE(Q)$. A detailed description of this process of determining $SPARSE(Q)$ is given in stages from “utilizing heavy jobs” to “sparsifying regions.” For an example of this process that uses the terminology later described in these stages, see Fig. 4. Any remaining time ranges not selected at this stage will effectively be dense time ranges, and are delegated into $RANGES(Q_L)$, $RANGES(Q_R)$ (after dealing with their alignment issues). This process of designating time ranges to delegate is detailed in stages from “creating dense ranges” to “resolving leaves.”

As a reminder, we have chosen $Z(Q)$ such that the total weight inside any grid slice (a time range between two consecutive endpoints of $Z(Q)$) of Q is at most $\frac{\varepsilon^2}{\log^2(N)}w(P(Q))$. Recall that $Z(Q)$ contains grid endpoints. For the heavy crossing job, consider the grid endpoint immediately to its left and to its right. Without loss of generality, consider the right one and call it r . How we proceed can be split into two cases:

1. In the first case, r overlaps a job J in OPT' with weight at most $\frac{\varepsilon}{\log(n)}w(OPT(Q))$. We delete J and draw a boundary at r . In doing this, we will charge the weight of J against the cell Q . There are at most two jobs we charge in this manner for that original heavy interval, one for the grid endpoint to the right and one to the left. Meaning, each cell will be charged in this manner at most twice for a total of $O(\frac{\varepsilon}{\log(n)}w(OPT))$ weight at each level and $O(\varepsilon w(OPT))$ weight overall.
2. In the other case, r overlaps a job J that has weight greater than $\frac{\varepsilon}{\log(n)}w(OPT(Q))$. We call J a *highlighted* job. Our algorithm proceeds by considering the grid endpoint immediately to the right of J . We determine what to do with this grid endpoint in a recursive manner. Meaning, we proceed in the same two cases that we did when considering what to do with r , and continue this recursive process until we finally draw a boundary.

After this process, we will have drawn a *region* (time range corresponding to where we drew a left and right boundary for) in which OPT' has the one heavy crossing job, a number of highlighted jobs (possibly zero), and potentially some remaining jobs that are neither crossing nor highlighted (we call these *useless*). It is our goal to convert this region into time ranges that we can use as sparse time ranges. Our process also guarantees this region has borders with endpoints in $Z(Q)$. Note that we have created a region within some time range of $RANGES(Q)$, but not every point in the time range is necessarily contained within the region.

Deleting useless jobs.

In the generated region, we define *useless* jobs as all jobs that are neither crossing nor highlighted. Useless jobs are completely contained within grid slices. We want to convert the region into sparse time ranges, but there may be many useless jobs that make the region very dense. Thus, we will delete all jobs in the region that are useless. By the process of generating the region, any such job is fully contained within a grid slice for which there is a heavy crossing job or highlighted job partially overlapping the grid slice. We charge deletion of all useless jobs in a given slice by charging against a highlighted or heavy crossing job that must partially overlap the given slice. By definition of $Z(Q)$, useless jobs in the slice add up to a total weight of at most $\frac{\epsilon^2}{\log^2(n)}w(OPT(Q))$. This is because we set $Z(Q)$ with $X = \frac{\log^2(n)}{\epsilon^2}$ and thus the optimal solution within any grid slice has total weight at most $\frac{\epsilon^2}{\log^2(n)}w(OPT(Q))$.

Moreover, $\frac{\epsilon^2}{\log^2(N)}w(OPT(Q))$ is at least a factor of ϵ less than the highlighted or heavy crossing job we are charging against (and there are only two such slices whose useless jobs are charging against any highlighted or heavy jobs).

Sparsifying the region.

Now, the region only contains heavy crossing job or highlighted jobs. We aim to split the region into ranges for $SPARSE(Q)$ without deleting much weight. The region may have more than $\frac{1}{\epsilon}$ jobs (meaning it is not sparse). If this is the case, we desire to split the region into time ranges that each have $\leq \frac{1}{\epsilon}$ jobs and start/end at grid endpoints of $Z(Q)$. To do so, we number the jobs in a region from left to right and consider them in groups based on their index modulo $\frac{1}{\epsilon}$. Note that a group does not consist of consecutive jobs. Then, we delete the group with lowest weight. We delete this group because we make the observation that all remaining jobs in the region must contain a grid endpoint within it. This is because heavy crossing jobs must contain a grid endpoint by how we defined $Z(Q)$, and highlighted jobs must contain a grid endpoint by their definition. Thus, we can delete the jobs belonging to the lightest group and split the time range at the grid endpoints contained inside each of the deleted jobs. In doing so, we lose at most a factor of ϵ of the total weight of all the considered jobs. However, now each resulting time range will have at most $\frac{1}{\epsilon}$ jobs and thus will be a valid sparse range in $SPARSE(Q)$ (because for any range containing a number of consecutive jobs greater than $\frac{1}{\epsilon}$, we will have split it). Note that all these sparse ranges have endpoints in $Z(Q)$. With all of its terminology now defined, readers may find the example illustrated in Fig. 4 helpful for their understanding.

Snapping dense ranges.

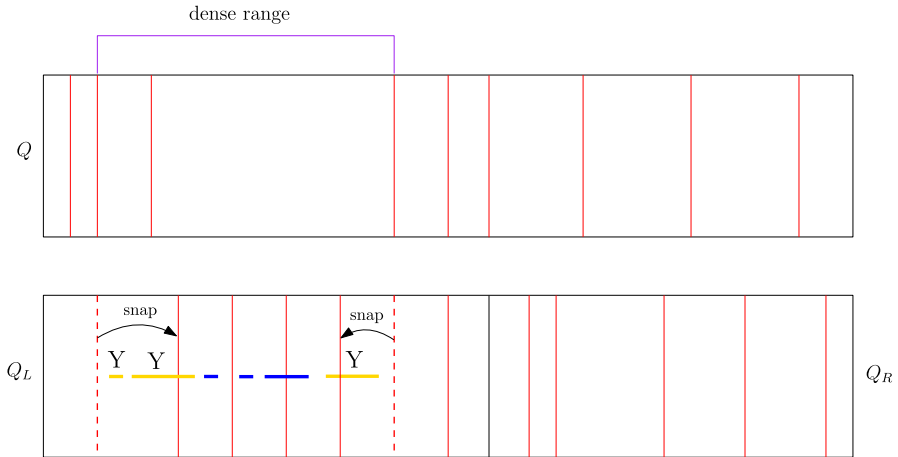


Fig. 5 This example illustrates why the snapping we perform has to be done with care. The horizontal segments in this figure represent jobs. We show an initial dense range (outlined in purple) with endpoints in $Z(Q)$. With dashed vertical lines, we show where these endpoints are in Q_L . Importantly, they are not aligned with $Z(Q_L)$, i.e., the vertical dashed lines do not belong to $Z(Q_L)$. However, our structure *requires* that dense ranges align with $Z(Q_{child})$, so we must address this. If we were to naively snap the endpoints of the dense range inward to the endpoints of $Z(Q_L)$, then we would need to delete some jobs (these deleted jobs are colored in yellow and marked by “Y”), while some other jobs would not be affected (like the remaining jobs in this example, those colored in blue). While this naive snapping may be fine in some cases, it will incur significant loss in cases in which the “Y” jobs have large weight. Notice that naively snapping outward to define a new region corresponding to the purple one is not a solution either, as this could cause the dense time range to overlap with a previously selected sparse time range. Having overlapping ranges can cause us to choose intersecting jobs, and thus an invalid solution. Thus, we detail a more comprehensive manner of dealing with snapping

Recall that not all of the time ranges that we are modifying from $RANGES(Q)$ were part of the region. In particular, there are the time ranges originally in $RANGES(Q)$ other than \mathcal{R}_{cross} , as well as the time range in \mathcal{R}_{cross} to the left of the region, and to the right of the region. We call these remaining time ranges our *dense ranges* because they may contain many jobs. Note how all dense range are now completely contained within Q_L or Q_R . Ideally, we assign dense ranges to $RANGES(Q_L)$ or $RANGES(Q_R)$. However, the remaining dense time ranges have one remaining potential issue, that their endpoints may not align with $Z(Q_{child})$ even though they align with $Z(Q)$. For an example of this issue, see Fig. 5. The core of this problem is that these dense time ranges correspond to time ranges we would like to delegate to children of Q (i.e., add to $RANGES(Q_L)$ and $RANGES(Q_R)$). However, there is the requirement that time ranges delegated to $RANGES(Q_L)$ and $RANGES(Q_R)$ must have endpoints in $Z(Q_L)$ and $Z(Q_R)$, respectively. Therefore, we have to modify the dense ranges so they align with the grid endpoints of one of Q 's children. It is tempting to naively “snap” the endpoints of these time ranges inward to the nearest grid endpoints of $Z(Q_{child})$, meaning to slightly contract the endpoints of the time ranges inward so they align with $Z(Q_{child})$. Unfortunately, this might result in some jobs being ignored in the process (as illustrated in Fig. 5); a cell does not consider jobs which are not within a given range. If these ignored jobs have non-negligible total reward, ignoring them can result

in a poor solution. In the stage “snapping dense ranges” we detail a more involved contraction-like snapping process that contracts inwards similar to our argument for expanding outwards from heavy crossing jobs when we determined sparse ranges. In our contraction-like snapping process, we convert some of the beginning and end of the dense range into sparse ranges, so we do not need to delete some of the high-reward jobs that we would need to delete with naive snapping. In the stages from “using essential jobs” to “resolving leafs”, we detail how to apply modifications to fulfill the required properties and how to analyze the contraction process with charging arguments.

Consider an arbitrary unaligned dense time range U . Ideally, we would “snap” the endpoints of U inward to the nearest grid point of $Z(Q_{child})$ (i.e. move the left endpoint of U to the closest grid point of $Z(Q_{child})$ to its right, and the right endpoint of U to the closest grid endpoint of $Z(Q_{child})$ to its left). However, doing so may force us to delete a job in OPT' that is too valuable (as we would have to delete jobs that overlap the section of U that was snapped inwards). So, we will handle U differently. Without loss of generality, suppose we want to “snap” inward the left endpoint of U to align with $Z(Q_{child})$. Doing so may leave some jobs outside the snapped range. We define the cost of snapping as the total weight of jobs that were previously contained within the range but are no longer completely contained within after snapping. If immediately snapping inward the left endpoint to the nearest grid point of $Z(Q_{child})$ would cost at most $\frac{2\varepsilon}{\log^2(n)} w(OPT(Q))$, we do that immediately. Otherwise, this snapping step would cost more than $\frac{2\varepsilon}{\log^2(n)} w(OPT(Q))$, implying that there is a job that overlaps with the grid endpoint of $Z(Q_{child})$ to the right of U 's left endpoint (all other jobs we are forced to delete are strictly inside a slice of $Z(Q_{child})$ and thus have total weight $\leq \frac{\varepsilon^2}{\log^2(n)} w(OPT(Q_{child})) \leq \frac{\varepsilon^2}{\log^2(n)} w(OPT(Q))$) and has weight of at least $\frac{2\varepsilon}{\log^2(n)} w(OPT(Q)) - \frac{\varepsilon^2}{\log^2(n)} w(OPT(Q)) \geq \frac{\varepsilon}{\log^2(n)} w(OPT(Q))$. We mark that job as “essential”.

Then, we look to the right of that essential job and examine the job that overlaps the next grid endpoint to the right in $Z(Q_{child})$. If this job has weight at most $\frac{2\varepsilon}{\log^2(n)} w(OPT(Q))$, we delete it and draw a boundary. Otherwise, we mark it as “essential” and continue (following the same process). When we are done, we have a prefix of the dense time range that contains some number of “essential” jobs and other jobs, and then a border at a grid endpoint of $Z(Q_{child})$. The final “snapping” where we deleted jobs to add the split point had cost $\leq \frac{2\varepsilon}{\log^2(n)} w(OPT(Q))$. In essence, these essential jobs are the collection of jobs that were too valuable for us to delete them when we were undergoing the snapping process.

Using essential jobs.

We will assume this dense time range had a snapping budget and charge the aforementioned final snapping cost to that. Now, we just need to find a way to use the time range prefix with the essential jobs. We delete all jobs that are not essential in this time range with a similar argument as earlier, that such a job is completely contained in a grid slice with total weight of jobs $\leq \frac{\varepsilon^2}{\log^2(N)} w(OPT(Q))$ which is at most a factor of ε of an essential job partially contained within the slice (and it is partially

contained within at most two slices). Then, we convert this time range of essential jobs (with potentially many such essential jobs) into sparse time ranges in the same way as done previously during the “sparsifying regions” step. We do so by grouping the jobs according their index modulo $\frac{1}{\varepsilon}$, deleting the group with least total weight, and drawing a border at the grid endpoint of $Z(Q_{child})$ contained within the deleted jobs. Again, by our process we know all such essential jobs must contain a grid endpoint. This creates sparse time ranges with endpoints in $Z(Q) \cup Z(Q_{child})$ and our dense time range has endpoints in $Z(Q_{child})$ so they are both valid.

Financing a snapping budget.

Finally, we need to show that we actually have a sufficient snapping budget. Consider our dense time ranges. We may adjust their endpoints in other scenarios, but we only split dense time ranges into more dense time ranges when they are a crossing range. As only one there is only one crossing range at every cell Q , if we give the newly created range a snapping budget of $O(\frac{\varepsilon}{\log(n)} w(OPT(Q)))$, then we do not lose more than $O(\varepsilon w(OPT))$ in total. We showed above that each dense range will use at most $O(\frac{\varepsilon}{\log^2(N)} w(OPT(Q)))$ of its snapping budget at each level, so it will will use $O(\frac{\varepsilon}{\log(n)} w(OPT(Q)))$ in total and stay within its allotted budget of $O(\frac{\varepsilon}{\log(n)} w(OPT(Q)))$ throughout.

Resolving leaves.

Finally, when we have a time range but it cannot be delegated to Q_{child} because Q_{child} does not exist, note there is only possibly room for one job in the range (as by definition of the decomposition of Q , no job starts or ends in this range). So we simply consider this range as part of $SPARSE(Q)$.

This concludes the proof by providing a way to convert OPT to a solution OPT' that obeys our structure and is a $(1 + \varepsilon)$ -approximation of OPT . \square

5.3 Efficiently Approximating Sparse Solutions

Now, we focus on designing an efficient algorithm for approximating the optimal solution in a sparse time range.

Lemma 5.1 *Given any contiguous time range \mathcal{R} and an integer K , consider an optimal solution $OPT(\mathcal{R}, K)$ in \mathcal{R} containing at most K jobs and ignoring jobs with weight less than $\varepsilon/n \cdot w_{max}$. Then, there is an algorithm that in \mathcal{R} finds a $(1 + \varepsilon)$ -approximate solution to $OPT(\mathcal{R}, K)$ in $O\left(\frac{K \log(n) \log^2(K/\varepsilon)}{\varepsilon^2}\right)$ time and with at most $O\left(\frac{K \log(K/\varepsilon)}{\varepsilon}\right)$ jobs.*

Proof To prove this claim, we use a dynamic programming approach where our state is the total weight of jobs selected so far. The dynamic programming table EARLIEST contains for state X , EARLIEST[X], the earliest/leftmost point in time for which the total weight of X is achieved. If we implement this dynamic programming directly, it would require space proportional to the value of solution (which equals the largest possible X). Our goal is to avoid this time/space dependence. To that end, we design an *approximate* dynamic program that requires only poly-logarithmic dependence on the value of an optimal solution. We derive the following technical tool to enable this:

Claim 5.6 *Let S be the set of all powers of $(1 + \varepsilon/K)$ not exceeding W , i.e., $S = \{(1 + \varepsilon/K)^i \mid 0 \leq i \leq \lfloor \log_{1+\varepsilon/K} W \rfloor\}$. Consider an algorithm that supports the addition of any K values (each being at least 1) where the sum of these K values is guaranteed to not exceed W . The values are added one by one. After each addition step, the algorithm maintains a running-total by rounding down the sum of the new value being added and the previous rounded running-total to the nearest value in S . Then, the final running-total of the algorithm is a $(1 + \varepsilon)$ approximation of the true sum of those K values.*

Proof Consider the sequence of K values and thus K additions. Let OPT denote the exact sum of the K values. Let SOL denote the running-total we achieve at the end of our additions. Finally, let CUR_i denote the running-total as we do these additions at the beginning of stage i , which must be in S at the end of every stage. We prove that $SOL \geq (1 - \varepsilon)OPT$ and thus SOL is a $(1 + \varepsilon)$ approximation of OPT . Initially, $CUR_0 = 0$. Each step, we add some value v_i to CUR_i . This new value $CUR'_i = CUR_i + v_i$. Then, we round CUR'_i to the nearest power of $(1 + \varepsilon/K)$ and denote this as CUR''_i . We call the amount we lose by rounding down the loss $\ell_i = CUR'_i - CUR''_i$. For the next stage, we denote $CUR_{i+1} = CUR''_i$. Note that

$$\frac{\ell_i}{OPT} \leq \frac{\ell_i}{SOL} \leq \frac{\ell_i}{CUR''_i} = \frac{CUR'_i - CUR''_i}{CUR''_i} \leq \frac{\varepsilon}{K}$$

or, otherwise, we would have rounded to a different power of $(1 + \varepsilon/K)$. Thus, $\ell_i \leq OPT(\frac{\varepsilon}{K})$. Note that $SOL = CUR_K$ and $CUR_K + \sum_i \ell_i = OPT$. As such,

$$\begin{aligned} SOL &= OPT - \sum_i \ell_i \geq OPT - K \left(OPT \left(\frac{\varepsilon}{K} \right) \right) = OPT - \varepsilon \cdot OPT \\ &= (1 - \varepsilon)OPT. \end{aligned}$$

□

Inspired by Claim 5.6, we now define a set of states S as follows. Our states will represent powers of $(1 + \varepsilon/K)$ from 1 to Kw , and hence

$$|S| = O \left(\frac{\log(Kw)}{\log(1 + \varepsilon/K)} \right) = O \left(\frac{K \log(Kw)}{\varepsilon} \right).$$

Using this, we create a set of states S which corresponds to powers of $(1 + \varepsilon/K)$ from 1 to Kw (and 0). We want to maintain for each of these states, approximately the smallest prefix with at most K jobs where we could get total weight approximately equal to $s \in S$. To do this, we loop over the states in increasing order of value. Suppose the current state corresponds to having approximate weight $s \in S$ and $\text{EARLIEST}[s]$ is the shortest prefix we have that has approximate weight s . Then we loop over all rounded weights $v \in \{(1 + \varepsilon)^i\}$. There are $O(\log(w)/\varepsilon)$ such v . For each v , set \mathcal{V} to be the value of $s + v$ rounded down to the nearest power of $(1 + \varepsilon/K)$. Then, if the earliest ending time of a job with rounded weight v that starts after $\text{EARLIEST}[s]$ is less than $\text{EARLIEST}[\mathcal{V}]$, we update $\text{EARLIEST}[\mathcal{V}]$ to that ending time. We can calculate

the earliest ending time of any job, with a particular rounded weight, starting after some specified time in $O(\log(n))$ time by maintaining a balanced binary search tree (as done in Sect. 4) for each of the $O(\log(w)/\varepsilon)$ rounded weights (to powers of $(1 + \varepsilon)$). This negligibly adds $O(\log(n))$ time to each update. In total, this solution runs in $O\left(\frac{K \log(n) \log(w) \log(Kw)}{\varepsilon^2}\right)$ time.

As we can ignore all jobs with weight less than ε/nw_{max} , then we can only focus on jobs with weights in $[\frac{\varepsilon w_{max}}{n}, w_{max}]$ and effectively modify w to be n/ε by dividing all weights by $\frac{\varepsilon w_{max}}{n}$. This enables us to use $w = O(n/\varepsilon)$ in the above runtime bound. As such, this algorithm runs in $O\left(\frac{K \log(n) \log(n/\varepsilon) \log(Kn/\varepsilon)}{\varepsilon^2}\right)$ time.

To show the algorithm's correctness, observe that since we always round down, we will not overestimate the cost. Moreover, we show with that any set of K additions will be within a factor of $(1 + \varepsilon)$ from its true value. \square

Corollary 5.7 *For our application, we let $K = \frac{1}{\varepsilon}$. As such, we have a $(1 + \varepsilon)$ -approximation algorithm of the minimum solution with at most $\frac{1}{\varepsilon}$ jobs that runs in time*

$$O\left(\frac{K \log(n) \log(n/\varepsilon) \log(Kn/\varepsilon)}{\varepsilon^2}\right) = O\left(\frac{\log(n) \log^2(n/\varepsilon)}{\varepsilon^3}\right) = O\left(\frac{\log^3(n)}{\varepsilon^3}\right).$$

5.4 Dynamically Maintaining $Z(Q)$ —Proof of Lemma 5.4

Now, we describe how to maintain $Z(Q)$, to intelligently subdivide the cells with guarantees as restated below:

Lemma 5.4 (Dynamically maintaining $Z(Q)$) *For any fixed positive integer X , it is possible to return a set $Z(Q)$ for any cell Q in the hierarchical decomposition in $O(X \cdot \log^3(n))$ query time. Moreover, the returned $Z(Q)$ will satisfy the following properties:*

- For every Q , the optimal solution within each grid slice of $Z(Q)$ is at most $O(w(OPT(Q))/X)$; as a reminder, we ignore jobs with weights less than $w(OPT(Q)) \cdot \varepsilon/n$.
- For every Q , $|Z(Q)| = O(X \cdot \log^2(n))$

Proof Suppose how all jobs are rounded down to powers of 2. Note how for a cell Q , let $w_{max}(Q)$ correspond to the reward of the job with the largest reward contained completely within Q . Clearly, $OPT(Q) \geq w_{max}(Q)$. Moreover, by discarding all jobs with weight less than $\varepsilon/n \cdot w_{max}(Q)$, we discard jobs with total weight at most $\varepsilon \cdot w_{max}(Q) \leq \varepsilon \cdot OPT(Q)$. Accordingly, we focus just on jobs with weights in range $[\varepsilon/n \cdot w_{max}(Q), w_{max}(Q)]$. As these weights have been rounded to powers of 2, there are only $\lceil \log(\frac{w_{max}(Q)}{\varepsilon/n w_{max}(Q)}) \rceil = O(\log(n/\varepsilon))$ distinct remaining weights. Moreover, we assume that $1/\varepsilon \leq n$, as otherwise we can obtain a better algorithm by simply rerunning the classical static algorithm for each update. Altogether, this implies that it suffices to consider $O(\log(n))$ distinct weights.

In our approach, we consider each distinct weight independently, enabling us to consider a $Z^i(Q)$ for only jobs with rounded weight 2^i . That is, $Z^i(Q)$ is computed

with respect to a set of jobs all having the same weight, which enables us to treat $Z^i(Q)$ computation as if it was performed for the unweighted variant. At the end, we let $Z(Q)$ to be the union over the $O(\log(n))$ different $Z^i(Q)$, giving us a $Z(Q)$ with our desired guarantees. This approach is particularly desirable, as we will show how for a particular fixed weight, i.e., we consider the unweighted variant, we can use ideas very similar to those discussed in Sect. 4 to obtain the $Z^i(Q)$. Expanding our scope, for each rounded weight 2^i , let us maintain a constant-factor approximation of the unweighted problem using the border-based algorithm of Theorem 1.1. In other words, we run the algorithm of Theorem 1.1 with $\varepsilon' = O(1)$ such that it has update time $O(\log(n))$ and maintains an $O(1)$ -approximation of the unweighted interval scheduling problem.

Consider SOL^i to be **the set of points** corresponding to the border-based $O(1)$ -approximation when only considering jobs of rounded weight 2^i . In particular, SOL^i contains all start/endpoints of the selected jobs by the approximation, as well as all borders. $SOL^i[L, R]$ contains all points in SOL^i within $[L, R]$. We define $OPT([L, R], i)$ as the optimal **number of jobs** one can schedule when only considering jobs with rounded weights 2^i when only considering jobs fully contained within $[L, R]$.

Claim 5.8 For all i, L, R , it holds that: $OPT([L, R], i) \leq |SOL^i[L, R]|$

Proof Recall that the border-based approximation algorithm maintains a set of borders and finds the optimal solution within each border chunk. The optimal solution within is calculated by using the greedy earliest-ending algorithm. In general, consider any job J . This job J must contain an endpoint of a job in the approximately chosen solution, or it must contain a border. If this were not the case, there are only two possibilities: (i) J is completely contained within a job chosen by the approximate solution, or (ii) J does not intersect any job chosen by the approximate solution. For case (i): this is impossible as the greedy earliest-ending algorithm would not have chosen the job that completely contains J . For case (ii): this is impossible because J could be added to the solution within the corresponding border chunk, and this is impossible because the solution within each border chunk is optimal. As each job J must contain a point of $SOL^i[L, R]$, it must hold that $OPT([L, R], i) \leq |SOL^i[L, R]|$. \square

Also note a similar bound in the opposing direction:

Claim 5.9 For all i, L, R it holds that: $|SOL^i[L, R]| \leq 3 \cdot OPT([L, R], i) + 3$

Proof From $SOL^i[L, R]$, ignore the at most two points corresponding to endpoints of jobs that are only partially within $[L, R]$, and ignore the first remaining point if it corresponds to a border (for a total of ignoring at most 3 points). Of the remaining points in $SOL^i[L, R]$, they all correspond to endpoints of jobs fully within $[L, R]$, or are a border following such a job. Note how the number of these jobs with points corresponding to them in $SOL^i[L, R]$ must be at most $OPT([L, R], i)$ by definition. Accordingly, we will charge the two points from each job (and its associated border if there is one) to a different job corresponding from $OPT([L, R], i)$, for a total of at most 3 points of $SOL^i[L, R]$ being charged per job in $OPT([L, R], i)$. \square

All SOL^i can be maintained with update time $O(\log(n))$ because we only update one unweighted $O(1)$ -approximation per job insertion or deletion. We compute each

$Z^i(Q)$ for a cell Q corresponding to time range $[L, R]$, by taking $O(X \log(n))$ quantiles of $SOL^i[L, R]$. Each of these $Z^i(Q)$ can be achieved with $O(X \log(n))$ walks down a balanced binary search tree, resulting in $O(X \log^2(n))$ time. We define $Z(Q)$ as the union of the $O(\log(n))$ different $Z^i(Q)$. In total, $Z(Q)$ is obtained in $O(X \log^3(n))$ time and $|Z(Q)| = O(X \log^2(n))$.

Finally, the optimal solution within any grid slice, ignoring jobs with weight less than $\varepsilon \cdot w(OPT(Q))$, is upper-bounded by the union of the independent optimal solutions for each rounded weight. Within each grid slice of any $Z^i(Q)$, the optimal solution of jobs using weight 2^i is upper-bounded by $O(\frac{2^i |SOL^i[l_Q, r_Q]|}{X \log(n)}) = O(\frac{2^i OPT([l_Q, r_Q], i)}{X \log(n)}) = O(\frac{w(OPT(Q))}{X \cdot \log(n)})$ following from Claim 5.8, taking $X \log(n)$ quantiles of $SOL^i[l_Q, r_Q]$ to form $Z^i(Q)$, and Claim 5.9. Accordingly, bounding over the $O(\log(n))$ different $Z^i(Q)$, the optimal solution within each grid slice is at most $O(w(OPT(Q))/X)$. \square

5.5 Combining All Ingredients—Proof of Theorem 1.2

Now, we put this all together to get a cohesive solution that efficiently calculates an approximately optimal solution of the desired structure. When we handle an insertion/deletion, we make an update to the corresponding balanced binary search tree T . Recall that we use a balanced binary search tree such as a red-black tree so that T has depth $O(\log(n))$ and there are $O(\log(n))$ rotations per update. For the $O(\log(n))$ cells Q corresponding to nodes in T affected by rotations, we will recompute aspects of Q such as $Z(Q)$ and all $DP(Q, S)$. For each such cell Q , we will compute a sparse solution corresponding to each segment formed by considering all pairs of grid endpoints $Z(Q) \cup Z(Q_L) \cup Z(Q_R)$ and a dense solution for each segment S formed by pairs of endpoints $Z(Q)$ denoted as $DP(Q, S)$.

To compute all sparse solutions, we use $O(|Z(Q) \cup Z(Q_L) \cup Z(Q_R)|^2)$ calls to our algorithm from Lemma 5.1 resulting in $O(|Z(Q) \cup Z(Q_L) \cup Z(Q_R)|^2 (\frac{\log^3(n)}{\varepsilon^3})) = O(\frac{\log^8(n)}{\varepsilon^4} (\frac{\log^3(n)}{\varepsilon^3})) = O(\frac{\log^{11}(n)}{\varepsilon^7})$ running time. To obtain this complexity, we use the upper-bound $|Z(Q)| = O(X \cdot \log^2(n))$ from Lemma 5.4 and the fact that we let $X = \frac{\log^2 n}{\varepsilon^2}$ in the beginning of Sect. 5.2.3.

To compute all $DP(Q, S)$, we build on the proof of Lemma 5.5. Namely, from the proof of Lemma 5.5 a $(1 + \varepsilon)$ -approximate solution is maintained by dividing S into sparse, i.e., $SPARSE(Q)$, and dense segments of Q_L and Q_R , i.e., $RANGES(Q_L)$ and $RANGES(Q_R)$. We update our data structure from bottom to top. Hence, when we update $DP(Q_L)$ and $DP(Q_R)$ it enables us to learn approximate optimal values gained from a set $RANGES(Q_L)$ and $RANGES(Q_R)$. Thus, to calculate $DP(Q, S)$ we consider an interval scheduling instance where jobs start at a grid endpoint of S and end at a grid endpoint of S . In this instance, jobs correspond to all the sparse segments of $Z(Q)$, $Z(Q_L)$, $Z(Q_R)$ and all the dense segments of $Z(Q_L)$, $Z(Q_R)$. We compute this dense segment answer for all dense segments $Z(Q)$ in $O(|Z(Q) \cup Z(Q_L) \cup Z(Q_R)|^3) = O(\frac{\log^{12}(n)}{\varepsilon^6})$ time, with a dynamic program where the state is the starting and ending point of a segment and the transition tries all potential grid

endpoints to split the range at (or just uses the interval from the start to the end). For each update, we update $O(\log(n))$ cells affected by rotations by recomputing the optimal sparse solutions for segments and the respective $DP(Q, S)$. Finally, at the beginning of each update, we use $O(\log(n))$ calls to our algorithm for computing $Z(Q)$ from Sect. 5.4 with $X = \frac{\log^2(n)}{\varepsilon^2}$ in $O(X \cdot \log^3(n))$ time for $O(\frac{\log^5(n)}{\varepsilon^2})$ time for each cell. As such, our total update time is $O\left(\log(n) \cdot \left(\frac{\log^{11}(n)}{\varepsilon^7} + \frac{\log^{12}(n)}{\varepsilon^6} + \frac{\log^6(n)}{\varepsilon^2}\right)\right) = O\left(\frac{\log^{12}(n)}{\varepsilon^7} + \frac{\log^{13}(n)}{\varepsilon^6}\right)$.

Author Contributions All authors contributed to developing and discussing ideas, and writing and reviewing the manuscript.

Declarations

Conflict of interest We have no conflict of interest to disclose.

References

1. Salot, P.: A survey of various scheduling algorithm in cloud computing environment. *Int. J. Res. Eng. Technol.* **2**(2), 131–135 (2013)
2. Sharma, R., Soni, V.K., Mishra, M.K., Bhuyan, P.: A survey of job scheduling and resource management in grid computing. *World Acad. Sci Eng. Technol.* **64**, 461–466 (2010)
3. Kolen, A.W., Lenstra, J.K., Papadimitriou, C.H., Spieksma, F.C.: Interval scheduling: a survey. *Naval Res. Logist.* **54**(5), 530–543 (2007)
4. Mingozzi, A., Boschetti, M.A., Ricciardelli, S., Bianco, L.: A set partitioning approach to the crew scheduling problem. *Oper. Res.* **47**(6), 873–888 (1999)
5. Hochbaum, D.S., Maass, W.: Approximation schemes for covering and packing problems in image processing and VLSI. *J. ACM* **32**(1), 130–136 (1985)
6. Agarwal, P.K., Van Kreveld, M., Suri, S.: Label placement by maximum independent set in rectangles. *Comput. Geom.* **11**(3–4), 209–218 (1998)
7. Verweij, B., Aardal, K.: An optimisation algorithm for maximum independent set with applications in map labelling. In: *European Symposium on Algorithms*, pp. 426–437. Springer (1999)
8. Bore, S., Cardinal, J., Iacono, J., Koumoutsos, G.: Dynamic geometric independent set (2020). arXiv preprint [arXiv:2007.08643](https://arxiv.org/abs/2007.08643)
9. Henzinger, M., Neumann, S., Wiese, A.: Dynamic approximate maximum independent set of intervals, hypercubes and hyperrectangles. In: *36th International Symposium on Computational Geometry (SoCG 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik (2020)
10. Gawrychowski, P., Pokorski, K.: Sublinear dynamic interval scheduling (on one or multiple machines) (2022). arXiv preprint [arXiv:2203.14310](https://arxiv.org/abs/2203.14310)
11. Gavruskin, A., Khoussainov, B., Kokho, M., Liu, J.: Dynamic interval scheduling for multiple machines. In: *International Symposium on Algorithms and Computation*, pp. 235–246. Springer (2014)
12. Gavruskin, A., Khoussainov, B., Kokho, M., Liu, J.: Dynamic algorithms for monotonic interval scheduling problem. *Theoret. Comput. Sci.* **562**, 227–242 (2015)
13. Frank, A.: Some polynomial algorithms for certain graphs and hypergraphs. In: *Proceedings of the 5th British Combinatorial Conference*. Utilitas Mathematica (1975)
14. Tardos, E., Kleinberg, J.: *Algorithm Design*. Pearson (2005)
15. Arkin, E.M., Silverberg, E.B.: Scheduling jobs with fixed start and end times. *Discrete Appl. Math.* **18**(1), 1–8 (1987)
16. Lenstra, J.K., Rinnooy Kan, A.: Complexity of scheduling under precedence constraints. *Oper. Res.* **26**(1), 22–35 (1978)
17. Correa, J.R., Schulz, A.S.: Single-machine scheduling with precedence constraints. *Math. Oper. Res.* **30**(4), 1005–1021 (2005)

18. Robert, J., Schabanel, N.: Non-clairvoyant scheduling with precedence constraints. In: Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 491–500 (2008)
19. Skutella, M., Uetz, M.: Stochastic machine scheduling with precedence constraints. *SIAM J. Comput.* **34**(4), 788–802 (2005)
20. Buttazzo, G.C., Bertogna, M., Yao, G.: Limited preemptive scheduling for real-time systems: a survey. *IEEE Trans. Ind. Inf.* **9**(1), 3–15 (2012)
21. Pinedo, M.: *Scheduling*, vol. 29. Springer, Berlin (2012)
22. Levey, E., Rothvoss, T.: A $(1 + \epsilon)$ -approximation for makespan scheduling with precedence constraints using LP hierarchies. *SIAM J. Comput.* **8**, 16–201 (2019)
23. Cardinal, J., Iacono, J., Koumoutsos, G.: Worst-case efficient dynamic geometric independent set. In: 29th Annual European Symposium on Algorithms (ESA 2021), vol. 204, p. 25. Schloss Dagstuhl-Leibniz-Zentrum für Informatik (2021)

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

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.