



# The preemptive resource allocation problem

Kanthi Sarpatwar<sup>1</sup> · Baruch Schieber<sup>2</sup> · Hadas Shachnai<sup>3</sup>

Accepted: 17 May 2023 / Published online: 21 June 2023

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2023

## Abstract

We revisit a classical scheduling model to incorporate modern trends in data center networks and cloud services. Addressing some key challenges in the allocation of shared resources to user requests (jobs) in such settings, we consider the following variants of the classic *resource allocation problem* (RAP). The input to our problems is a set  $J$  of jobs and a set  $M$  of homogeneous hosts, each has an available amount of some resource. Assuming that time is slotted, a job is associated with a release time, a due date, a weight and a given length, as well as its resource requirement. A *feasible* schedule is an allocation of the resource to a subset of the jobs, satisfying the job release times/due dates as well as the resource constraints. A crucial distinction between classic RAP and our problems is that we allow preemption and migration of jobs, motivated by virtualization techniques. We consider two natural objectives: *throughput maximization* (MaxT), which seeks a maximum weight subset of the jobs that can be feasibly scheduled on the hosts in  $M$ , and *resource minimization* (MinR), that is finding the minimum number of (homogeneous) hosts needed to feasibly schedule all jobs. Both problems are known to be NP-hard. We first present an  $\Omega(1)$ -approximation algorithm for MaxT instances where time-windows form a laminar family of intervals. We then extend the algorithm to handle instances with arbitrary time-windows, assuming there is sufficient slack for each job to be completed. For MinR we study a more general setting with  $d$  resources and derive an  $O(\log d)$ -approximation for any fixed  $d \geq 1$ , under the assumption that time-windows are not too small. This assumption can be removed leading to a slightly worse ratio of  $O(\log d \log^* T)$ , where  $T$  is the maximum due date of any job.

**Keywords** Machine scheduling · Resource allocation · Vector packing · Approximation algorithms

## 1 Introduction

The proliferation of virtualization and containerization technologies, along with the advent of increasingly powerful multi-core processors, has made it possible to execute mul-

tiples virtual machines (or *jobs*) simultaneously on the same host, as well as to preempt and migrate jobs with relative ease. We address some fundamental problems in the efficient allocation of shared resources such as CPU cores, RAM, or network bandwidth to several competing jobs. These problems are modeled to exploit multi-job execution and facilitate preemption and migration, while respecting resource and timing constraints. Typically, the infrastructure service providers are oversubscribed; therefore, the common goals here include admission control of jobs to maximize throughput, or minimizing the additional resource required to process all jobs.

The broad setting considered in this paper is the following. Suppose we are given a set of jobs  $J$  that need to be scheduled on a set of identical hosts  $M$ , where each host has a limited amount of one or more resources. Each job  $j \in J$  has release time  $r_j$ , due date  $d_j$ , and length  $p_j$ , along with a required amount of the resource  $s_j$  ( $\bar{s}_j$  for multiple resources). We assume that time is slotted. A job  $j$  can be preempted and migrated across hosts but cannot be processed

---

This work was partly carried out during a visit of the third author to DIMACS with support from the Special Focus on Bridging Continuous and Discrete Optimization.

---

✉ Kanthi Sarpatwar  
sarpatwa@us.ibm.com  
Baruch Schieber  
baruch.m.schieber@njit.edu  
Hadas Shachnai  
hadas@cs.technion.ac.il

<sup>1</sup> IBM T. J. Watson Research Center, Yorktown Heights, NY, USA

<sup>2</sup> Computer Science Department, New Jersey Institute of Technology, Newark, NJ, USA

<sup>3</sup> Computer Science Department, Technion, Haifa, Israel

*simultaneously* on multiple hosts, i.e., at any time slot a job can be processed by at most one host. However, multiple jobs can be processed by any host at any given time, as long as their combined resource requirement does not exceed the available resource. We consider two commonly occurring objectives, namely, *throughput maximization* and *resource minimization*.

In the *maximum throughput* (MaxT) variant, we are given a set of homogeneous hosts  $M$  and a set of jobs  $J$ , such that each job  $j$  has a profit  $w_j > 0$  and attributes  $(p_j, s_j, r_j, d_j)$ . The goal is to find a subset  $S \subseteq J$  of jobs of maximum profit  $\sum_{j \in S} w_j$  that can be feasibly scheduled on  $M$ . This problem can be viewed as a preemptive variant of the classic *resource allocation problem* (RAP) (Phillips et al., 2000; Chen et al., 2002; Călinescu et al., 2011; Bansal et al., 2014).

In the *resource minimization* (MinR) variant, we assume that each job  $j$  has a resource requirement vector  $\bar{s}_j \in [0, 1]^d$  as one of the attributes, where  $d \geq 1$  is the number of available resources. W.l.o.g., we assume that each host has a unit amount of each of the  $d$  resources. A schedule that assigns a set of jobs  $S_{i,t}$  to a host  $i \in M$  at time  $t$  is feasible if  $\sum_{j \in S_{i,t}} \bar{s}_j \leq \bar{1}^d$ . Given a set of jobs  $J$  with attributes  $(p_j, \bar{s}_j, r_j, d_j)$ , we seek a set of (homogeneous) hosts  $M$  of minimum cardinality such that all of the jobs can be scheduled feasibly on  $M$ . MinR is a generalization of the classic *vector packing* (VP) problem, in which a set of  $d$ -dimensional items needs to be feasibly packed into a minimum number of  $d$ -dimensional bins of unit size in each dimension, i.e., the vector sum of all the items packed into each bin has to be less than or equal to  $\bar{1}^d$ . Any instance of VP can be viewed as an instance of MinR with  $r_j = 0$ ,  $d_j = 1$  and  $p_j = 1$  for job  $j \in J$ .

Another application of this general scheduling scenario relates to the allocation of space and time to advertisements by online advertisement platforms (such as Google or Facebook). In the *ad placement problem* (Dawande et al., 2003; Freund & Naor, 2004) we are given a schedule length of  $T$  time slots and a collection of ads that need to be scheduled within this time frame. The ads must be placed in a rectangular display area, whose contents can change in different time slots. All ads share the same height, which is the height of the display area, but may have different widths. Several ads may be displayed simultaneously (side by side), as long as their combined width does not exceed the width of the display area. In addition, each ad specifies a display count (in the range  $1, \dots, T$ ), which is the number of time slots during which the ad must be displayed. The actual time slots in which the advertisement will be displayed may be chosen arbitrarily by the scheduler, and, in particular, need not be consecutive. Suppose that each ad is associated with some positive profit, and the scheduler may accept or reject any given ad. A common objective is to schedule a maximum-profit subset of ads within a display area of given width.

Indeed, this problem can be cast as a special case of MaxT with a single host, where all jobs have the same release time and due date.

## 1.1 Prior work

Given an algorithm  $\mathcal{A}$  for a maximization (minimization) problem  $\Pi$ , let  $\mathcal{A}(I)$ ,  $OPT(I)$  denote the value of the solution output by  $\mathcal{A}$  and by an optimal solution for a problem instance  $I$ , respectively. For  $\rho \in (0, 1]$  ( $\rho \geq 1$ ), we say that  $\mathcal{A}$  is a  $\rho$ -approximation algorithm if, for any instance  $I$ ,  $\frac{\mathcal{A}(I)}{OPT(I)} \geq \rho$  ( $\frac{\mathcal{A}(I)}{OPT(I)} \leq \rho$ ).

The classical problem of preemptively scheduling a set of jobs with attributes  $(p_j, s_j = 1, r_j, d_j)$  on a single machine so as to maximize throughput can be cast as a special case of MaxT with a single host, where each job requires all of the available resource. Lawler (1990) showed that in this special case MaxT admits a *polynomial time approximation scheme* (PTAS), and the problem is polynomially solvable for uniform job weights. For multiple hosts (i.e.,  $m = |M| > 1$ ), this special case of MaxT ( $s_j = 1$  for all  $j \in J$ ) admits a  $\frac{1}{6+\varepsilon}$ -approximation, for any fixed  $\varepsilon > 0$ . This follows from a result of Kalyanasundaram and Pruhs (2001).

As mentioned earlier, another special case of MaxT was studied in the context of advertisement placement. The *ad placement* problem was introduced by Adler et al. (2002) and later studied in numerous papers (see, e.g., Dawande et al. (2003, 2005), Freund and Naor (2004), Kumar et al. (2007), Kaul et al. (2017) and the survey in Pandey et al. (2017)). Freund and Naor (2004) presented a  $(1/3 - \varepsilon)$ -approximation for the maximum profit version, namely, for MaxT with a single host and the same release time and due date for all jobs.

Fox and Korupolu (2013) studied our preemptive scheduling model, with job attributes  $(p_j, s_j, r_j, d_j)$ , under another popular objective, namely, minimizing weighted flow-time. Their work differs from ours in two ways: while they focus on the online versions, we consider our problems in an offline setting. Further, as they note, while the throughput and resource minimization objectives are also commonly considered metrics, their techniques only deal with flow-time. In fact, these objectives are fundamentally different, and we need novel algorithms to tackle them.

The non-preemptive variant of MaxT, known as the *resource allocation problem* (RAP), was introduced by Phillips et al. (2000), and later studied by many authors (see, e.g., Bar-Noy et al. (2001b, a); Călinescu et al. (2011); Chakaravarthy et al. (2014); Jain et al. (2015); Chen et al. (2002) and the references therein).<sup>1</sup> Chakaravarthy et al. (2014) consider a generalization of RAP and obtain a constant approximation based on a primal-dual algorithm. We

<sup>1</sup> RAP is also known as the *bandwidth allocation problem*.

note that the preemptive versus non-preemptive problems differ quite a bit in their structural properties.

As mentioned above, MinR generalizes the classic vector packing (VP) problem. The first non-trivial  $O(\log d)$ -approximation algorithm for VP was presented by Chekuri and Khanna (2004), for any fixed  $d \geq 1$ . This ratio was improved by Bansal et al. (2009) to a randomized algorithm with asymptotic approximation ratio arbitrarily close to  $\ln d + 1$ . Bansal et al. (2016) recently improved this ratio further to  $0.807 + \ln(d + 1)$ . A “fractional variant” of MinR was considered by Jansen and Porkolab (2002), where time was assumed to be continuous. For this problem, in the case of a single host, they obtain a PTAS, by solving a configuration linear program (rounding the LP solution is not necessary because time is continuous in their case).

Resource minimization was considered also in the context of the ad placement problem. In this variant, all ads must be scheduled, and the objective is to minimize the width of the display area required to make this possible. Freund and Naor (2004) gave a 2-approximation algorithm for the problem, which was later improved by Dawande et al. (2005) to  $3/2$ . This implies a 3-approximation for MinR instances with  $d = 1$ , where all jobs have the same release time and due date. We believe that this ratio can be slightly improved, using the property that  $s_j \leq 1$  for all  $j \in J$ . Indeed, we can schedule the jobs to use the resource, such that the total resource requirements at any two time slots differ at most by one. Thus, the total amount of resource required at any time exceeds the optimum,  $OPT$ , at most by one unit, implying the jobs can be feasibly scheduled on  $2OPT + 1$  hosts.

Another line of work relates to the non-preemptive version of MinR, where  $d = 1$  and the requirement of each job is equal to 1 (see, e.g. (Chuzhoy et al., 2004; Chuzhoy & Codenotti, 2009)); thus, at most one job can be scheduled on each host at any time.

### 1.2 Contributions and techniques

Denote the time window for processing job  $j \in J$  by  $\chi_j = [r_j, d_j]$ , and let  $|\chi_j| = d_j - r_j + 1$  denote the length of the interval. Assume that  $\min_j r_j = 1$ . Let  $T = \max_j d_j$ . Without loss of generality assume that every time slot  $t \in [T]$  belongs to a time window of some job. Otherwise, we can “break” the problem into smaller sub-problems at the time slots  $t \in [T]$  that do not belong to any time window. Let  $r = \arg \max_j |\chi_j|$ . Assume also that  $\sum_{j \in J} p_j > |\chi_r|$ . This assumption is valid since otherwise job  $r$  is guaranteed to be scheduled on a single host after the rest of the jobs are scheduled. Thus, we can remove job  $r$  and consider the remaining jobs.

Since the output needs to specify the times each job is executed, the output size is  $\Omega(\sum_{j \in J} p_j)$ , and certainly  $\Omega(n)$ . We note that our algorithms are polynomial in  $T$  (and in

$n$ ). It follows from our assumptions that  $|\chi_r|$  and thus also  $T = O(n|\chi_r|)$  are polynomial in the output size.

For summarizing our results, we need the notion of *slackness*. Throughout the discussion, we assume that time windows are large enough, i.e., there is a constant  $\lambda \in (0, 1)$ , such that  $p_j \leq \lambda|\chi_j|$  for any job  $j$ . Such an assumption is quite reasonable in scenarios arising in our applications. We call  $\lambda$  the *slackness* parameter of the instance.

For the MaxT problem, we present (in Sect. 3) an  $\Omega(1)$ -approximation algorithm. As mentioned above, the non-preemptive version of this problem is the classic RAP. To see the structural differences between the non-preemptive and preemptive versions, we consider their natural linear programming relaxations. In solving RAP it suffices to have a variable  $x_{jt}$  for each job  $j$  and time slot  $t$ , indicating the start of job  $j$  at slot  $t$ . This allows to apply a natural randomized rounding algorithm, where job  $j$  is scheduled to start at time  $t$  with probability  $x_{jt}$ . On the other hand, in MaxT a job can be preempted; therefore, each job requires multiple indicator variables. Further, these variables must be rounded in an *all-or-nothing* fashion, i.e., either we schedule all parts of a job or none of them. Our approach to handle this situation is to, somewhat counter-intuitively, “dumb down” the linear program by not committing the jobs to a particular schedule; instead, we choose a subset of jobs that satisfy certain knapsack constraints and construct the actual schedule in a subsequent phase.

We first consider a *laminar* variant of the problem, where the time windows for the jobs are chosen from a laminar family of intervals.<sup>2</sup> This setting includes several important special cases, such as (i) all jobs are released at  $t = 0$  but have different due dates, or (ii) jobs are released at different times, but all must be completed by a given due date. Recall that  $m = |M|$  is the number of hosts. Our result for the laminar case is a  $\frac{1}{2} - \lambda(\frac{1}{2} + \frac{1}{m})$ -approximation algorithm, assuming that the slackness parameter satisfies  $\lambda < 1 - \frac{2}{m+2}$ . Using a simple transformation of an arbitrary instance to laminar, we obtain a  $\frac{1}{8} - \lambda(\frac{1}{2} + \frac{1}{m})$ -approximation algorithm for general instances, assuming that  $\lambda < \frac{1}{4} - \frac{1}{2(m+2)}$ . Our results imply that as  $\lambda$  decreases, the approximation ratio approaches  $\frac{1}{2}$  and  $\frac{1}{8}$  for the laminar and the general case, respectively.

Subsequently, we tighten the slackness assumption further to obtain an  $\Omega(1)$ -approximation algorithm for any constant slackness  $\lambda \in (0, 1)$  for the laminar case and any constant  $\lambda \in (0, \frac{1}{4})$  for the general case. In the special case where the weight of the job is equal to its area,<sup>3</sup> we extend an algorithm due to Chen et al. (2002) to obtain an  $\Omega(1)$ -approximation ratio for the general case with no assumption on slackness.

<sup>2</sup> See the formal definition in Sect. 2.

<sup>3</sup> The *area* of  $j \in J$  is the total resource requirement of  $j$  (see Sect. 2.)

Our algorithm for the laminar case relies on a non-trivial combination of a *packing* phase and a *scheduling* phase. While the first phase ensures that the output solution has high profit, the second phase guarantees its feasibility. To facilitate a successful completion of the selected jobs, we formulate a set of conditions that must be satisfied in the packing phase. Both phases make use of the structural properties of a *laminar* family of intervals. In the packing phase, we apply our rounding procedure (for the LP solution) to the tree representation of the intervals.<sup>4</sup> We further use this tree in the scheduling phase, to feasibly assign the resource to the selected jobs in a bottom-up fashion. Our framework for solving MaxT is general, and may therefore find use in other settings of *non-consecutive* resource allocation.

For the MinR problem, we obtain (in Sect. 5) an  $O(\log d)$ -approximation algorithm for any constant  $d \geq 1$ , under a mild assumption that any job has a window of size  $\Omega(d^2 \log d \log T)$ . We show that this assumption can be removed, leading to a slight degradation in the approximation factor to  $O(\log d \log^* T)$ , where  $\log^* T$  is the smallest integer  $\kappa$  such that  $\underbrace{\log \log \dots \log T}_{\kappa \text{ times}} \leq 1$ . Our approach builds on

a formulation of the problem as a configuration LP, inspired by the works of Bansal et al. (2009), Fleischer et al. (2011). However, we quickly deviate from these prior approaches, in order to handle the time windows and the extra constraints. Our algorithm involves two main phases: a *maximization phase* and *residual phase*. Roughly speaking, a configuration is a subset of jobs that can be feasibly assigned to a host at a given time slot  $t$ .

Initially, we solve a configuration LP, in which we find the minimum number of hosts,  $m$ , needed to schedule all jobs, and a set of configurations that are (fractionally) selected at any time  $t \in [T]$ . For each  $t$ , we then choose  $O(m \log d)$  configurations with probabilities proportional to their LP-values. In this phase, jobs may be allocated the resource only for part of their processing length. In the second phase, we construct a residual instance based on the amount of time each job has been processed. A key challenge is to show that, for any time window  $\chi$ , the total “area” of jobs left to be scheduled is at most  $1/d$  of the original total area. We use this property to solve the residual instance.

## 2 Preliminaries

We start with some definitions and notation. For our preemptive variants of RAP, we assume w.l.o.g. that each host has a unit amount of each resource. We further assume that time is slotted. We allow non-consecutive allocation of a resource

<sup>4</sup> This procedure bears some similarity to the *pipage* rounding technique of Ageev and Sviridenko (2004).

to each job, as well as job migration. Multiple jobs can be assigned to the same machine at a given time, but no job can be processed by multiple machines at the same time. Formally, we denote the set of jobs assigned to host  $i \in M$  at time  $t$  by  $S_{i,t}$ . We say that job  $j$  is *completed* if there are  $p_j$  time slots  $t \in [r_j, d_j] = \chi_j$  in which  $j$  is allocated its required amount of the resource on some host. A job  $j$  is completed if  $|\{t \in \chi_j : \exists i \in M \text{ such that } j \in S_{i,t}\}| \geq p_j$ . Let  $T = \max_{j \in J} d_j$  be the latest due date of any job.

In MaxT, each job  $j \in J$  has a resource requirement  $s_j \in (0, 1]$ . An assignment of a subset of jobs  $S \subseteq J$  to the hosts in  $M$  is feasible if each job  $j \in S$  is completed, and for any time slot  $t$  and host  $i \in M$ ,  $\sum_{j \in S_{i,t}} s_j \leq 1$ , i.e., the sum of requirements of all jobs assigned to host  $i$  is at most the available resource. For the MinR variant, we assume multiple resources. Thus, each job  $j$  has a resource requirement vector  $\bar{s}_j \in [0, 1]^d$ , for some constant  $d \geq 1$ . Further, each host has a unit amount of each of the  $d$  resources. An assignment of a set of jobs  $S_{i,t}$  to a host  $i \in M$  at time  $t$  is feasible if  $\sum_{j \in S_{i,t}} \bar{s}_j \leq \bar{1}^d$ .

Let  $a_j = s_j p_j$  denote the total resource requirement (or, *area*) of job  $j \in J$  and refer to the quantity  $w_j/a_j$  as the *density* of job  $j$ . Finally, a set of intervals is *laminar* if for any two intervals  $\chi'$  and  $\chi''$ , exactly one of the following holds:  $\chi' \subseteq \chi''$ ,  $\chi'' \subseteq \chi'$  or  $\chi' \cap \chi'' = \emptyset$ .

## 3 Throughput maximization

We first consider the case where  $\mathcal{L} = \{\chi_j : j \in J\}$  forms a *laminar* family of intervals. In Sect. 3.1, we present an  $\Omega(1)$ -approximation algorithm for the laminar case when  $\lambda \in (0, 1 - \frac{2}{m+2})$ . Following this, we describe in Sect. 3.2 our constant approximation for the general case for  $\lambda \in (0, \frac{1}{4} - \frac{1}{2(m+2)})$ . We then show, in Sect. 3.3, how to tighten the results to any constant slackness parameter (i)  $\lambda \in (0, 1)$  in the laminar case (ii)  $\lambda \in (0, \frac{1}{4})$  in the general case. As an interesting corollary, we obtain an  $\Omega(\frac{1}{\log n})$ -approximation algorithm for the general MaxT problem with no slackness assumption. Later, in Sect. 4, we show that in the special case of maximum utilization (i.e., the profit of each job equals its “area”), we obtain an  $\Omega(1)$  guarantee with no assumption on the slackness.

### 3.1 The Laminar case

Our algorithm proceeds in two phases. While the first phase ensures that the output solution has high profit, the second phase guarantees its feasibility. Specifically, let  $\omega \in (0, 1 - \frac{\lambda}{m}]$  be a parameter (to be determined).



In Phase 1, we find a subset of jobs  $S$  satisfying a knapsack constraint for each time-window  $\chi$ . Indeed, any feasible solution guarantees that the total area of jobs within any time-window  $\chi \in \mathcal{L}$  is at most  $m|\chi|$ . Our knapsack constraints further restrict the total area of jobs in  $\chi$  to some fraction of  $m|\chi|$ . We adopt an LP-rounding based approach to compute a subset  $S$  that is optimal subject to the further restricted knapsack constraints. (We remark that a dynamic programming approach would work as well. However, such an approach would not provide us with any intuition as to how an optimal solution for the further restricted instance compares with the optimal solution of the original instance.)

In Phase 2 we allocate the resource to the jobs in  $S$ , by considering separately each host  $i$  at a given time slot  $t \in [T]$  as a unit-sized bin  $(i, t)$  and iteratively assigning each job  $j \in S$  to a subset of such available bins, until  $j$  has the resource allocated for  $p_j$  distinct time slots. An outline of the two phases is given in Algorithm 1.

**Input:** Set of jobs  $J$ , hosts  $M$  and a parameter  $\omega \in (0, 1 - \frac{\lambda}{m}]$   
**Output:** Subset of jobs  $S \subseteq J$  and a feasible assignment of  $S$  to the hosts in  $M$   
**Phase 1:** Select a subset  $S \subseteq J$ , such that for each  $\chi \in \mathcal{L}$ :  
 $\sum_{j \in S: \chi_j \subseteq \chi} a_j \leq (\omega + \frac{\lambda}{m})m|\chi|$   
**Phase 2:** Find a feasible allocation of the resource to the jobs in  $S$

**Algorithm 1:** Throughput maximization algorithm outline

**Phase 1** The algorithm starts by finding a subset of jobs  $S \subseteq J$  such that for any  $\chi \in \mathcal{L}$ :  $\sum_{j \in S: \chi_j \subseteq \chi} a_j \leq (\omega + \frac{\lambda}{m})m|\chi|$ . We solve the following LP relaxation, in which we impose stricter constraint on the total area of the jobs assigned in each time window  $\chi$ .

$$\begin{aligned} \text{LP : Maximize } & \sum_{j \in J} w_j x_j \\ \text{Subject to: } & \sum_{j: \chi_j \subseteq \chi} a_j x_j \leq \omega m |\chi| \quad \forall \chi \in \mathcal{L} \\ & 0 \leq x_j \leq 1 \quad \forall j \in J \end{aligned}$$

**Rounding the Fractional Solution** Suppose  $\mathbf{x}^* = (x_j^* : j \in J)$  is an optimal fractional solution for the LP. Our goal is to construct an integral solution  $\hat{\mathbf{x}} = (\hat{x}_j : j \in J)$ . We refer to a job  $j$  with  $x_j^* \in (0, 1)$  as a *fractional job*, and to the quantity  $a_j x_j^*$  as its *fractional area*. W.l.o.g., we may assume that for any interval  $\chi \in \mathcal{L}$ , there is at most one job  $j$  with  $\chi_j = \chi$  such that  $0 < x_j^* < 1$ , i.e., it is fractional. Indeed, if two such jobs exist, then the fractional value of the higher density job (breaking ties arbitrarily) can be increased to obtain a solution no worse than the optimal. Note, however, that there could be fractional jobs  $j'$  with  $\chi_{j'} \subset \chi$ .

We start by setting  $\hat{x}_j = x_j^*$  for all  $j \in J$ . Consider the tree representation of  $\mathcal{L}$ , which contains a node (also denoted by  $\chi$ ) for each  $\chi \in \mathcal{L}$ , and an edge between nodes corresponding to  $\chi$  and  $\chi'$ , where  $\chi' \subset \chi$ , if there is no interval  $\chi'' \in \mathcal{L}$

such that  $\chi' \subset \chi'' \subset \chi$ .<sup>5</sup> Our rounding procedure works in a *bottom-up* fashion. As part of this procedure, we label the nodes with one of two possible colors: *gray* and *black*. Initially, all leaf nodes are colored *black*, and all internal nodes are colored *gray*. The procedure terminates when all nodes are colored *black*. A node  $\chi$  is colored as *black* if the following property holds:

**Property 1** For any path  $\mathcal{P}(\chi, \chi_l)$  from  $\chi$  to a leaf  $\chi_l$  there is at most one fractional job  $j$  such that  $\chi_j$  lies on  $\mathcal{P}(\chi, \chi_l)$ .

We note that the property trivially holds for the leaf nodes. Now, consider a *gray* interval  $\chi$  with children  $\chi_1, \chi_2, \dots, \chi_v$ , each colored *black*. Note that  $\chi$  is well defined because leaf intervals are all colored black. If there is no fractional job that has  $\chi$  as its time-window, Property 1 follows by induction, and we color  $\chi$  *black*. Assume now that  $j$  is a fractional job that has  $\chi$  as its time-window (i.e.,  $\chi_j = \chi$ ). If there is no other fractional job that has its time-window (strictly) contained in  $\chi$ , Property 1 is trivially satisfied. Therefore, assume that there are other fractional jobs  $j_1, j_2, \dots, j_l$  that have their time-windows (strictly) contained in  $\chi$ . Now, we decrease the fractional area (i.e., the quantity  $a_j \hat{x}_j$ ) of  $j$  by  $\Delta$  and increase the fractional area of jobs in the set  $\{j_1, j_2, \dots, j_l\}$  by  $\Delta_k$  for job  $j_k$ , such that  $\Delta = \sum_{k \in [l]} \Delta_k$ . Formally, we set  $\hat{x}_j \rightarrow \hat{x}_j - \frac{\Delta}{a_j}$  and  $\hat{x}_{j_k} \rightarrow \hat{x}_{j_k} + \frac{\Delta_k}{a_{j_k}}$ . We choose these increments  $\Delta_k$  such that either  $\hat{x}_j$  becomes 0, or for each  $k \in [l]$ ,  $\hat{x}_{j_k}$  becomes 1. Clearly, in both scenarios, Property 1 is satisfied, and we color  $\chi$  *black*.

When all nodes are colored *black*, we round up the remaining fractional jobs. Namely, for all jobs  $j$  such that  $\hat{x}_j \in (0, 1)$ , we set  $\hat{x}_j = 1$ . It is important to note that by doing so we may violate the knapsack constraints. However, in Theorem 1, we bound the violation.

**Theorem 1** Given a slackness parameter  $\lambda \in (0, 1)$ , let  $\mathcal{I} = (J, M, \mathcal{L})$  be a laminar instance of MaxT with optimal profit  $W$  and  $\forall j \in J: p_j \leq \lambda|\chi_j|$ . For any  $\omega \in (0, 1 - \frac{\lambda}{m}]$ , the subset  $S = \{j \in J : \hat{x}_j = 1\}$ , obtained as above, satisfies  $\sum_{j \in S} w_j \geq \omega W$ , and for any  $\chi \in \mathcal{L}$ ,  $\sum_{j \in S: \chi_j \subseteq \chi} a_j \leq (\omega + \frac{s_{\max} \lambda}{m})m|\chi|$ , where  $s_{\max} = \max_j s_j$ .

**Proof** We first observe that any optimal solution  $\mathbf{x}^*$  for the LP satisfies:  $\sum_{j \in J} w_j x_j^* \geq \omega W$ . Indeed, consider an optimal solution  $O$  for the instance  $\mathcal{I}$ . We can construct a fractional feasible solution  $\mathbf{x}'$  for the LP by setting  $x'_j = \omega$  if  $j \in O$ ; otherwise,  $x'_j = 0$ . Clearly,  $\mathbf{x}'$  is a feasible solution for the LP with profit  $\omega W$ .

Consider an integral solution  $\hat{\mathbf{x}}$ , obtained by applying the rounding procedure on  $\mathbf{x}^*$ . We first show that  $\sum_{j \in J} w_j \hat{x}_j \geq$

<sup>5</sup> Throughout the discussion we use interchangeably the terms *node* and *interval* when referring to a time-window  $\chi \in \mathcal{L}$ .

$\omega W$ . To this end, we prove that  $\sum_{j \in J} w_j \hat{x}_j \geq \sum_{j \in J} w_j x_j^* \geq \omega W$ . Suppose we decrease the fractional area of a job  $j$  by an amount  $\Delta$ , i.e., we set  $\hat{x}_j \leftarrow \hat{x}_j - \frac{\Delta}{a_j}$ . By the virtue of our procedure, we must simultaneously increase the fractional area of some subset of jobs  $F_j$ , where for each  $k \in F_j$  we have  $\chi_k \subset \chi_j$ . Further, the combined increase in the fractional area of the jobs in  $F_j$  is the same  $\Delta$ . Now, we observe that the density of job  $j$  (i.e.,  $\frac{w_j}{a_j}$ ) cannot be higher than any of the jobs in  $F_j$ . Indeed, if  $j' \in F_j$  has density strictly lower than  $j$ , then the optimal solution  $\mathbf{x}^*$  can be improved by decreasing the fractional area of  $j'$  by some  $\epsilon > 0$  while increasing that of  $j$  by the same amount (it is easy to see that no constraint is violated in this process) – a contradiction. Therefore, our rounding procedure will never result in a loss, and  $\sum_{j \in J} w_j \hat{x}_j \geq \sum_{j \in J} w_j x_j^* \geq \omega W$ .

Let  $s_{\max} = \max_j s_j$ . We now show that, for each  $\chi \in \mathcal{L}$ ,  $\sum_{j \in J: \chi_j \subseteq \chi} a_j \hat{x}_j \leq (\omega + \frac{\lambda}{m})m|\chi|$ . First, observe that for any gray interval  $\chi$  the total fractional area is *conserved*. This is true because there is no positive transfer of fractional area to the subtree rooted at  $\chi$  from a node outside this subtree until  $\chi$  is colored black. Now, consider an interval  $\chi$  that is colored *black*. We note that for any job  $j$  with  $x_j^* = 0$ , our algorithm ensures that  $\hat{x}_j = 0$ , i.e., it creates no new fractional jobs.

Consider the vector  $\hat{\mathbf{x}}$  when the interval  $\chi$  is converted from *gray* to *black*. At this stage, we have that the total (fractional) area packed in the subtree rooted at  $\chi$  is  $V(\chi) \stackrel{\text{def}}{=} \sum_{j \in J: \chi_j \subseteq \chi} a_j \hat{x}_j \leq \omega m|\chi|$ . Let  $\mathcal{F}(\chi)$  denote the set of all fractional jobs  $j'$  that have their time-windows contained in  $\chi$  (i.e.,  $\chi_{j'} \subseteq \chi$ ). We claim that the maximum increase in  $V(\chi)$  by the end of the rounding procedure is at most  $\sum_{j' \in \mathcal{F}(\chi)} a_{j'}$ . This holds since our procedure does not change the variables  $\hat{x}_j \in \{0, 1\}$ . Thus, the maximum increase in the total area occurs due to rounding all fractional jobs into complete ones, after all nodes are colored black.

To complete the proof, we now show that the total area of the fractional jobs in the subtree rooted at  $\chi$  satisfies  $\mathcal{A}[\chi] \stackrel{\text{def}}{=} \sum_{j' \in \mathcal{F}(\chi)} a_{j'} \leq s_{\max} \lambda |\chi|$ . We prove this by induction on the level of node  $\chi$ . Clearly, if  $\chi$  is a leaf then the claim holds, since there can exist at most one fractional job  $j$  in  $\chi$ , and  $a_j = s_j p_j \leq s_{\max} \lambda |\chi|$ . Suppose that  $\{\chi_1, \chi_2, \dots, \chi_l\}$  are the children of  $\chi$ . If there is a fractional job  $j$  with  $\chi_j = \chi$  then, by Property 1, there are no other fractional jobs with time-windows contained in  $\chi$ . Hence,  $\mathcal{A}[\chi] = a_j \leq s_{\max} \lambda |\chi|$ . Suppose there is no fractional job with  $\chi_j = \chi$ ; then, by the induction hypothesis:  $\mathcal{A}[\chi_k] \leq s_{\max} \lambda |\chi_k|$  for all  $k \in [l]$ . Further,  $\sum_{k \in [l]} |\chi_k| \leq |\chi|$  and  $\mathcal{A}[\chi] = \sum_{k \in [l]} \mathcal{A}[\chi_k] \leq \sum_{k \in [l]} s_{\max} \lambda |\chi_k| \leq s_{\max} \lambda |\chi|$ .  $\square$

Let  $O$  be an optimal solution for  $\mathcal{I}$  satisfying:  $\forall \chi \in \mathcal{L} : \sum_{j \in O: \chi_j \subseteq \chi} a_j \leq cm|\chi|$ , for some  $c \geq 1$ . Then it is easy to verify that any optimal solution  $\mathbf{x}^*$  for the LP sat-

isfies:  $\sum_{j \in J} w_j x_j^* \geq \frac{\omega}{c} W$ . The next result follows from the proof of Theorem 1.

**Corollary 1** *Suppose  $\mathcal{I} = (J, M, \mathcal{L})$  is a laminar instance of  $\text{MaxT}$ , such that  $\forall j \in J : p_j \leq \lambda |\chi_j|$ . Let  $S^+ \subseteq J$  be a subset of jobs of total profit  $W$  satisfying  $\forall \chi \in \mathcal{L} : \sum_{j \in S^+: \chi_j \subseteq \chi} a_j \leq cm|\chi|$ , for some  $c \geq 1$ . Then, for any  $\omega \in (0, 1 - \frac{\lambda}{m}]$ , there exists a subset  $S \subseteq J$  satisfying  $\sum_{j \in S} w_j \geq \frac{\omega}{c} W$ , such that  $\forall \chi \in \mathcal{L}, \sum_{j \in S: \chi_j \subseteq \chi} a_j \leq (\omega + \frac{\lambda}{m})m|\chi|$ .*

**Phase 2** In Phase 1 we obtained a subset  $S \subseteq J$ , such that for each  $\chi \in \mathcal{L} : \sum_{j \in S: \chi_j \subseteq \chi} a_j \leq (\omega + \frac{\lambda}{m})m|\chi|$ . We now show that it is always possible to find a feasible packing of all jobs in  $S$ . We refer to host  $i$  at time  $t$  as a *bin*  $(i, t)$ . In the allocation phase we label a *bin* with one of three possible colors: *white*, *gray* or *black*. Initially, all bins are colored *white*. We color a bin  $(i, t)$  gray when some job  $j$  is assigned to host  $i$  at time  $t$  and color it black when we decide to assign no more jobs to this bin. Our algorithm works in a bottom-up fashion and marks an interval  $\chi$  as *done* when it has successfully completed all the jobs  $j$  with  $\chi_j \subseteq \chi$ . Consider an interval  $\chi$  such that any  $\chi' \subset \chi$  has already been marked *done*. Let  $j \in S$  be a job with time-window  $\chi_j = \chi$ , that has not been processed yet. To complete job  $j$ , we must pick  $p_j$  distinct time slots in  $\chi$  and assign it to a bin in each slot. Suppose that we have already assigned the job to  $p'_j < p_j$  slots so far. Denote by  $\text{avail}(j) \subseteq \chi$  the subset of bins in time slots where  $j$  has not been assigned yet. We pick the next slot and bin as shown in Algorithm 2.

```

1: if there exists a gray bin  $(i, t)$  in  $\text{avail}(j)$  then
2:   let  $S_{(i,t)}$  be the set of jobs assigned to this bin
3:   if  $\sum_{j' \in S_{(i,t)}} s_{j'} + s_j \leq 1$  then
4:     assign  $j$  to host  $i$  at time  $t$ 
5:   else if there exists a white bin  $(i', t')$  in  $\text{avail}(j)$  then
6:     assign  $j$  to host  $i'$  at time  $t'$ .
7:     color  $(i, t)$  and  $(i', t')$  black
8:     pair up  $(i, t) \leftrightarrow (i', t')$ 
9:   else
10:    report fail
11:  end if
12: else if there exists a white bin  $(i, t)$  in  $\text{avail}(j)$  then
13:   assign  $j$  to host  $i$  at time  $t$ 
14:   color the bin  $(i, t)$  gray
15: else
16:   report fail
17: end if

```

**Algorithm 2:** Resource allocation to job  $j$  in a single time slot

**Theorem 2** *For any  $\lambda < 1 - \frac{2}{m+2}$ , there is a  $\frac{1}{2} - \lambda (\frac{1}{2} + \frac{1}{m})$ -approximation algorithm for the laminar  $\text{MaxT}$  problem, assuming that  $p_j \leq \lambda |\chi_j|$  for all  $j \in J$ .*

**Proof** Given an instance  $\mathcal{I} = (J, M, \mathcal{L})$  and a parameter  $\omega \in (0, 1 - \frac{\lambda}{m}]$ , let  $W$  denote the optimal profit. We apply Theorem 1 to find a subset of jobs  $S \subseteq J$  of profit  $\omega W$ , such that for any  $\chi \in \mathcal{L}$ :  $\sum_{j \in S: \chi_j \subseteq \chi} a_j \leq (\omega + \frac{\lambda}{m})m|\chi|$ . We now show that there is a feasible resource assignment to the jobs in  $S$  for  $\omega = \frac{1}{2} - \lambda(\frac{1}{2} + \frac{1}{m})$ . Clearly, this would imply the theorem.

We show that for the above value of  $\omega$  Algorithm 2 never reports **fail**, i.e., the resource is feasibly allocated to all jobs in  $S$ . Observe that whenever bins  $(i, t)$  and  $(i', t')$  are paired up in Algorithm 2 and colored black  $\sum_{j \in S(i,t)} s_j + \sum_{j' \in S(i',t')} s_{j'} > 1$ . Thus, the total number of black bins  $< 2(\omega + \frac{\lambda}{m})m|\chi|$ . Also, a time slot  $t$  can contain at most one gray bin (that is, at most one bin  $(i, t)$  is gray). This is true since whenever a bin  $(i, t) \in \text{avail}(j)$  becomes gray there are no existing gray bins in  $\text{avail}(j)$ .

Assume towards a contradiction that Algorithm 2 reports **fail** while assigning job  $j$ . We say that time slot  $t \in \chi = \chi_j$  is *bad* if either time slot  $t$  contains a gray bin, or  $j$  is already assigned to some bin  $(i, t)$ . We first show that the following invariant holds. As long as no job  $j^+$  such that  $\chi \subset \chi_{j^+}$  has been allocated the resource: the number of bad time slots while processing job  $j$  is at most  $\lambda|\chi|$ . To show that the invariant holds just before job  $j$  is allocated the resource we distinguish between two cases. (i) If job  $j$  is the first job with time window  $\chi$  that is allocated the resource, then assuming that the invariant holds in each of the child intervals of  $\chi$ ,  $\{\chi_1, \chi_2 \dots, \chi_l\}$ , before any job with time window  $\chi$  is allocated the resource, implies that the number of bad time slots just before job  $j$  is allocated the resource is at most  $\sum_{k \in [l]} \lambda|\chi_k| \leq \lambda|\chi|$ . (ii) Job  $j$  is not the first job with time window  $\chi$  that is allocated the resource; then, assuming the invariant holds after the previous job with time window  $\chi$  is allocated the resource implies that it holds just before job  $j$  is allocated the resource.

Now, consider the iteration in which Algorithm 2 assigns  $j$  to host  $i$  at time slot  $t$ . If time slot  $t$  was bad before the assignment, then the number of bad time slots remains the same. Suppose that time slot  $t$  was not bad before the assignment. Clearly, it becomes bad as  $j$  is assigned to bin  $(i, t)$ . In this case, bin  $(i, t)$  must have been white before job  $j$  is assigned and it is either black or gray after the assignment. (a) If bin  $(i, t)$  is black after the assignment then Algorithm 2 must have considered a gray bin  $(i_1, t_1)$  and failed to assign  $j$  to host  $i_1$  at time  $t_1$ . Consequently, bin  $(i_1, t_1)$  is paired with  $(i, t)$  and also colored black. At this point time slot  $t_1$  which was bad before (since it contained the gray bin  $(i_1, t_1)$ ) is not bad any more. Thus, the number of bad time slots remains the same. (b) If bin  $(i, t)$  is gray after the assignment then there were no gray bins in  $\text{avail}(j)$  before and after this assignment, and thus the number of bad time slots is at most  $p_j \leq \lambda|\chi|$ .

It follows that the total number of bins that are either black or in bad time slots is  $< (\lambda + 2(\omega + \frac{\lambda}{m}))m|\chi|$ . Now, setting  $\omega = \frac{1}{2} - \lambda(\frac{1}{2} + \frac{1}{m})$ , there should be at least one bin  $(i^*, t^*)$  that is neither black nor in a bad time slot. But in this case, bin  $(i^*, t^*)$  must be white and in  $\text{avail}(j)$ , and Algorithm 2 could have assigned  $j$  to host  $i^*$  at time  $t^*$  – a contradiction to the assumption that the algorithm reports a **fail**.  $\square$

For convenience, we restate the claim shown in the proof of Theorem 2.

**Corollary 2** *Let  $\mathcal{I} = (J, M, \mathcal{L})$  be a laminar instance where  $p_j \leq \lambda|\chi_j| \forall j \in J$ , for  $\lambda \in (0, 1 - \frac{2}{m+2})$ . Let  $S \subseteq J$  be a subset of jobs, such that for any  $\chi \in \mathcal{L}$ :  $\sum_{j \in S: \chi_j \subseteq \chi} a_j \leq (\omega + \frac{\lambda}{m})m|\chi|$ , where  $\omega \leq \frac{1}{2} - \lambda(\frac{1}{2} + \frac{1}{m})$ . Then, there exists a feasible resource assignment to the jobs in  $S$ .*

### 3.2 The general case

We use a simple transformation of general instances of MaxT into laminar instances and prove an  $\Omega(1)$ -approximation ratio. Let  $\mathcal{W}$  denote the set of all time-windows for jobs in  $J$ , i.e.,  $\mathcal{W} = \{\chi_j : j \in J\}$ . We now construct a laminar set of intervals  $\mathcal{L}$  and a mapping  $\mathfrak{L} : \mathcal{W} \rightarrow \mathcal{L}$ . Recall that  $T = \max_{j \in J} d_j$ . The construction is done via a binary tree  $\mathcal{T}$  whose nodes correspond to intervals  $[l, r] \subseteq [T]$ . The construction is described in Algorithm 3.

**Input:** Job set  $J$  and  $\mathcal{W} = \{\chi_j : j \in J\}$   
**Output:** Laminar set of intervals  $\mathcal{L}$  and a mapping  $\mathfrak{L} : \mathcal{W} \rightarrow \mathcal{L}$

- 1: let  $[T]$  be the root node of tree  $\mathcal{T}$
- 2: **while**  $\exists$  a leaf node  $[l, r]$  in  $\mathcal{T}$  such that  $r - l > 1$  **do**
- 3:     add to  $\mathcal{T}$  two nodes  $[l, \lfloor \frac{l+r}{2} \rfloor]$  and  $[\lfloor \frac{l+r}{2} \rfloor + 1, r]$  as the children of  $[l, r]$
- 4: **end while**
- 5: let  $\mathcal{L}$  be the set of intervals corresponding to the nodes of  $\mathcal{T}$
- 6: For each  $\chi \in \mathcal{W}$ , let  $\mathfrak{L}(\chi) = \chi'$ , where  $\chi'$  is the largest interval in  $\mathcal{L}$  contained in  $\chi$ , breaking ties by picking the *rightmost* interval.

**Algorithm 3:** Transformation into a laminar set

**Lemma 1** *In Algorithm 3, the following property holds. For  $\chi \in \mathcal{L}$ , let  $\tilde{\chi} = \{t \in \chi_j : j \in J, \mathfrak{L}(\chi_j) = \chi\}$ , i.e., the union of all time-windows in  $\mathcal{W}$  that are mapped to  $\chi$ . Then,  $|\tilde{\chi}| \leq 4|\chi|$ .*

**Proof** For  $\chi \in \mathcal{L}$ , consider first  $j \in J$  such that  $\mathfrak{L}(\chi_j) = \chi$ . We use in the proof the next claim.  $\square$

**Property 1**  $\chi_j$  cannot completely contain 3 consecutive intervals in  $\mathcal{L}$  that are at the same level as  $\mathfrak{L}(\chi_j)$ .

**Proof** Suppose  $\chi_j$  contains at least 3 such consecutive intervals. Then, by the virtue of our algorithm,  $\mathfrak{L}(\chi_j)$  is the

rightmost interval. Let  $\hat{\chi}$  be the parent of  $\mathcal{L}(\chi_j)$ . Two cases arise:

*Case 1*  $\mathcal{L}(\chi_j)$  is a left child of  $\hat{\chi}$ . Consider the two other consecutive intervals at the same level as  $\mathcal{L}(\chi_j)$  that are contained in  $\chi_j$ . Observe that these two intervals are siblings; therefore, their parent (which is also in  $\mathcal{L}$ ) is also contained in  $\chi_j$ . This is a contradiction to the assumption that  $\mathcal{L}(\chi_j)$  is the largest interval in  $\mathcal{L}$  contained in  $\chi_j$ .

*Case 2*  $\mathcal{L}(\chi_j)$  is a right child of  $\hat{\chi}$ . We observe that the sibling of  $\mathcal{L}(\chi_j)$  must also be contained in  $\chi_j$ , implying that  $\hat{\chi}$  is contained in  $\chi_j$ , a contradiction.  $\square$

Now, for any  $\chi = [s, d] \in \mathcal{L}$ , let  $\chi_l = [s_l, d_l] \in \mathcal{W}$  (resp.  $\chi_r = [s_r, d_r]$ ) be the leftmost (resp. rightmost) interval in  $\mathcal{W}$  such that  $\mathcal{L}(\chi_l) = \chi$  (resp.  $\mathcal{L}(\chi_r) = \chi$ ); then,  $\tilde{\chi} = [s_l, d_r]$ . Consider the intervals  $\chi_1 = [s_l, s]$  and  $\chi_2 = [d, d_r]$ . By Claim 1,  $\chi_l$  cannot contain 3 consecutive intervals in  $\mathcal{L}$  at the same level as  $\chi$ . Thus,  $|\chi_1| < 2|\chi|$ . Also,  $|\chi_2| < |\chi|$ ; otherwise, there is an interval to the right of  $\chi$  of the same size that can be mapped to  $\chi_r$ . Thus,  $|\chi_1| + |\chi_2| < 3|\chi|$ . Now, the claim follows by observing that  $|\tilde{\chi}| = |\chi_1| + |\chi| + |\chi_2| \leq 4|\chi|$ .  $\square$

**Theorem 3** For any  $\lambda < \frac{1}{4} - \frac{1}{2(m+2)}$ , there exists a  $\frac{1}{8} - \lambda(\frac{1}{2} + \frac{1}{m})$ -approximation algorithm for MaxT, assuming that  $p_j \leq \lambda|\chi_j|$  for all  $j \in J$ .

**Proof** Given an instance  $(J, M, \mathcal{W})$  of MaxT with slackness parameter  $\lambda \in (0, 1)$ , we first use Algorithm 3 to obtain a laminar set of intervals  $\mathcal{L}$  and the corresponding mapping  $\mathcal{L} : \mathcal{W} \rightarrow \mathcal{L}$ . Consider a new laminar instance  $(J_\ell = \{j_\ell : j \in J\}, M_\ell = M, \mathcal{L})$ , constructed by setting  $\chi_{j_\ell} = \mathcal{L}(\chi_j)$ . Note that if  $S_\ell \subseteq J_\ell$  is a feasible solution for this new instance, the corresponding set  $S = \{j : j_\ell \in S_\ell\}$  is a feasible solution for the original instance. Let  $\lambda_\ell$  denote the slackness parameter for the new instance. We claim that  $\lambda_\ell \leq 4\lambda$ . Assume this is not true, i.e., there exists a job  $j_\ell$ , such that  $p_{j_\ell} > 4\lambda|\chi_{j_\ell}|$ ; however, by Lemma 1, we have  $p_{j_\ell} = p_j \leq \lambda|\chi_j| \leq 4\lambda|\chi_{j_\ell}|$ . A contradiction. Now, suppose  $O \subseteq J$  is an optimal solution of total profit  $W$  for the original (non-laminar) instance. Consider the corresponding subset of jobs  $O_\ell = \{j_\ell : j \in O\}$ . By Lemma 1, for any  $\chi \in \mathcal{L}$ ,  $|\tilde{\chi}| \leq 4|\chi|$ . It follows that, for any  $\chi \in \mathcal{L}$ ,  $\sum_{j_\ell \in O_\ell : \chi_{j_\ell} \subseteq \chi} a_{j_\ell} = \sum_{j \in O : \mathcal{L}(\chi_j) \subseteq \chi} a_j \leq 4m|\chi|$ . Now, we use Corollary 1 for the laminar instance, taking  $c = 4$ ,  $S^+ = O_\ell$  and  $\lambda_\ell \in (0, 1 - \frac{2}{m+2})$ . Then, for any  $\omega \in (0, 1 - \frac{\lambda_\ell}{m}]$ , there exists  $S_\ell \subseteq J_\ell$  of total profit  $\sum_{j_\ell \in S_\ell} w_{j_\ell} \geq \frac{\omega}{c}W$ , such that  $\forall \chi \in \mathcal{L}$ ,  $\sum_{j_\ell \in S_\ell : \chi_{j_\ell} \subseteq \chi} a_{j_\ell} \leq (\omega + \frac{\lambda_\ell}{m})m|\chi|$ . By Corollary 2, there is a feasible assignment of the resource to the jobs in  $S_\ell$  for  $\omega \leq \frac{1}{2} - \lambda_\ell(\frac{1}{2} + \frac{1}{m})$ . Taking  $\omega = \frac{1}{2} - 4\lambda(\frac{1}{2} + \frac{1}{m}) \leq \frac{1}{2} - \lambda_\ell(\frac{1}{2} + \frac{1}{m})$ , we have the approximation ratio  $\frac{w}{c} = \frac{1}{8} - 4\lambda(\frac{1}{8} + \frac{1}{4m})$ , for any

$\lambda < \frac{1}{4} - \frac{1}{2(m+2)}$ . We now return to the original instance and take for the solution the set  $S = \{j : j_\ell \in S_\ell\}$ .  $\square$

### 3.3 Relaxing the slackness requirements

In this section we show that the slackness requirements in Theorems 2 and 3 can be relaxed, while maintaining a constant approximation ratio for MaxT. In particular, for laminar instances, we show below that Algorithm 1 can be used to obtain a polynomial time  $\Omega(1)$ -approximation for any constant slackness parameter  $\lambda \in (0, 1)$ . For general MaxT instances, this leads to an  $\Omega(1)$ -approximation for any constant  $\lambda \in (0, \frac{1}{4})$ . We also show a polynomial time  $\Omega(\frac{1}{\log n})$ -approximation algorithm for general MaxT using no assumption on slackness. We use below the next result, for instances with ‘large’ resource requirement.

**Lemma 2** For any  $\delta \in (0, 1)$  there is an  $\Omega(\frac{1}{\log(1/\delta)})$ -approximation for any instance  $\mathcal{I} = (J, M, \mathcal{W})$  of MaxT satisfying  $s_j \geq \delta \forall j \in J$ .

**Proof** Given an instance  $\mathcal{I}$ , we first round down the resource requirement (or, height) of each job  $j \in J$  to the nearest value of the form  $\delta(1 + \varepsilon')^k$ , for some fixed  $\varepsilon' \in (0, 1)$  and integer  $0 \leq k \leq \lceil \log_{1+\varepsilon'}(\frac{1}{\delta}) \rceil$ . We now partition the jobs into  $O(\log(\frac{1}{\delta}))$  classes, such that the jobs in each class have the same rounded height. For a class with job height  $\delta(1 + \varepsilon')^k$ , let  $m_k = m \cdot \lfloor \frac{1}{\delta(1+\varepsilon')^k} \rfloor$ . We define for this class the instance  $\mathcal{I}_k = (J_k, M_k, \mathcal{W})$  of MaxT in which  $|M_k| = m_k$  and  $s_j = 1$  for all  $j \in J_k$ .

Lawler (1990) gave a PTAS for MaxT on a single host, where  $s_j = 1$  for all  $j \in J$ . Consider an algorithm  $\mathcal{A}_k$  for MaxT on  $\mathcal{I}_k$ , which proceeds as follows. We schedule iteratively the jobs in  $J$  on hosts  $1, \dots, m_k$ . Let  $\mathcal{J}_{i-1}$  be the set of jobs scheduled on hosts  $1, \dots, i-1$ , and  $\mathcal{J}_0 = \emptyset$ . In iteration  $i \geq 1$ , we use the PTAS of Lawler (1990) for the set of jobs  $J \setminus \mathcal{J}_{i-1}$ . We note that the resulting schedule uses no migrations. By a result of Kalyanasundaram and Pruhs (2001), this iterative algorithm (called in Kalyanasundaram and Pruhs (2001) Repeat-A) yields a ratio of  $\frac{1}{6+\varepsilon}$  to the profit of an optimal schedule for MaxT (which may use migrations).<sup>6</sup>

Let  $\mathcal{A}_k(\mathcal{I}_k)$  be the profit of the solution obtained for  $\mathcal{I}_k$ . Then we choose the solution set for the instance  $\mathcal{I}_{\ell^*}$  which maximizes the profit. That is,  $\mathcal{A}_{\ell^*}(\mathcal{I}_{\ell^*}) = \max_{0 \leq k \leq \lceil \log_{1+\varepsilon'}(\frac{1}{\delta}) \rceil} \mathcal{A}_k(\mathcal{I}_k)$ . We note that since the job heights are rounded down, transforming back to the original job heights requires the algorithm to make changes that we describe next. W.l.o.g., assume that  $m_{\ell^*} > m$  (otherwise, the rounded height of the scheduled jobs is larger than  $\frac{1}{2}$ , implying they can be scheduled feasibly with their original

<sup>6</sup> See Theorem 3.1 in Kalyanasundaram and Pruhs (2001).



heights on  $m$  hosts). Thus, among the  $m_{\ell^*}$  hosts, we select  $m'_{\ell^*} = \lfloor \frac{m_{\ell^*}}{1+\varepsilon'} \rfloor$  hosts on which the total weight of scheduled jobs is maximized. It follows that the approximation ratio is at least  $(\frac{1}{1+\varepsilon'} - \frac{1}{m_{\ell^*}}) \cdot \frac{1}{(6+\varepsilon)\lceil \log_{1+\varepsilon'}(\frac{1}{\delta}) \rceil}$ .  $\square$

### 3.3.1 Laminar instances

Recall that  $m = |M|$  is the number of hosts. Given a fixed  $\lambda \in (0, 1)$ , let

$$\alpha = \alpha(m, \lambda) = \frac{\lambda(1 - \lambda)}{1 - \lambda + \frac{\lambda}{m}}. \tag{1}$$

We note that  $\alpha < \lambda$ . In Phase 1 of Algorithm 1, we round the LP solution to obtain a subset of jobs  $S \subseteq J$ . We first prove the following.

**Lemma 3** *Let  $\lambda \in (0, 1)$  be a slackness parameter, and*

$$\omega = (1 - \alpha)(1 - \lambda) - \frac{\alpha\lambda}{m}, \tag{2}$$

where  $\alpha$  is defined in (1). Then, given a laminar instance  $\mathcal{I} = (J, M, \mathcal{L})$  satisfying  $p_j \leq \lambda|\chi_j|$  and  $s_j \leq \alpha$ , there is a feasible allocation of the resource to the jobs in the set  $S$  selected in Phase 1 of Algorithm 1.

**Proof** We generate a feasible schedule of the jobs in  $S$  proceeding bottom-up in each laminar tree. That is, we start handling job  $j$  only once all the jobs  $\ell$  with time windows  $\chi_\ell \subset \chi_j$  have been scheduled. Jobs having the same time window are scheduled in an arbitrary order. Let  $j$  be the next job, whose time window is  $\chi = \chi_j$ . We can view the interval  $\chi$  as a set of  $|\chi|$  time slots, each consisting of  $m$  unit size bins. We say that a time slot  $t \in \chi_j$  is ‘bad’ for job  $j$  if there is no space for one processing unit of  $j$  (i.e., an ‘item’ of size  $s_j$ ) in any of the bins in  $t$ ; else, time slot  $t$  is ‘good’. We claim that immediately before we start scheduling job  $j$  the number of bad time slots for  $j$  is at most  $\frac{m|\chi|(1-\lambda)(1-\alpha)-a_j}{m(1-s_j)}$ . Indeed, by Theorem 1, choosing for  $\omega$  the value in (2), after rounding the LP solution, the total area of jobs  $\ell \in S$  such that  $\chi_\ell \subseteq \chi_j$  is at most

$$(\omega + \frac{s_{max}\lambda}{m})m|\chi| \leq ((1 - \alpha)(1 - \lambda) - \frac{\alpha\lambda}{m} + \frac{\alpha\lambda}{m})m|\chi|. \tag{3}$$

Excluding job  $j$ , we have that the total occupied area in  $\chi$  is at most  $m|\chi|(1 - \lambda)(1 - \alpha) - a_j$ . In addition, for a time slot  $t$  to be ‘bad’ for job  $j$ , each bin in  $t$  has to be at least  $(1 - s_j)$ -full. This shows our claim.

Hence, the number of good time slots for  $j$  is at least

$$\begin{aligned} |\chi| - \frac{m|\chi|(1 - \lambda)(1 - \alpha) - a_j}{m(1 - s_j)} \\ &= |\chi|(1 - \frac{(1 - \lambda)(1 - \alpha)}{1 - s_j}) + \frac{a_j}{m(1 - s_j)} \\ &\geq \frac{p_j}{\lambda}(1 - \frac{(1 - \lambda)(1 - \alpha)}{1 - s_j}) + \frac{a_j}{m(1 - s_j)} \geq p_j \end{aligned}$$

The first inequality follows from the fact that  $p_j \leq \lambda|\chi_j| = \lambda|\chi|$ , and the second inequality holds since  $s_j \leq \alpha$ . Hence, job  $j$  can be feasibly scheduled, for any  $j \in S$ .  $\square$

Using Lemmas 2 and 3, we prove our main result.

**Theorem 4** *For any  $m \geq 1$  and constant  $\lambda \in (0, 1)$ , MaxT admits a polynomial time  $\Omega(1)$ -approximation on any laminar instance  $\mathcal{I} = (J, M, \mathcal{L})$  with slackness parameter  $\lambda$ .*

**Proof** Given a laminar instance  $\mathcal{I}$  satisfying the slackness condition, we handle separately two subsets of jobs.

**Subset 1** Jobs  $j$  satisfying  $s_j \leq \alpha = \alpha(m, \lambda)$ , where  $\alpha$  is defined in (1). We solve MaxT for these jobs using Algorithm 1, taking the value of  $\omega$  as in (2). By Theorem 1, the approximation ratio is  $\omega = (1 - \alpha)(1 - \lambda) - \frac{\alpha\lambda}{m} = (1 - \lambda)^2$ , i.e., we have a constant factor.

**Subset 2** For jobs  $j$  satisfying  $s_j > \alpha$ , use Lemma 2 to obtain an  $\Omega(\frac{1}{\log(1/\alpha)})$ -approximation.

Taking the best among the solutions for the two subsets of jobs, we obtain an  $\Omega(1)$ -approximation.  $\square$

### 3.3.2 The general case

Recall that, given a general MaxT instance,  $(J, M, \mathcal{W})$ , with a slackness parameter  $\lambda \in (0, 1)$ , our transformation yields a new laminar instance  $(J_\ell = \{j_\ell : j \in J\}, M_\ell = M, \mathcal{L})$  with a slackness parameter  $\lambda_\ell \leq 4\lambda$  (see the proof of Theorem 3). Now, define

$$\alpha_\ell = \alpha_\ell(m, \lambda_\ell) = \frac{\lambda_\ell(1 - \lambda_\ell)}{1 - \lambda_\ell + \frac{\lambda_\ell}{m}}, \tag{4}$$

and set

$$\omega = (1 - \alpha_\ell)(1 - \lambda_\ell) - \frac{\alpha_\ell\lambda_\ell}{m}. \tag{5}$$

Then, by Lemma 3, we have that any job  $j_\ell \in J_\ell$  selected for the solution set  $S$  can be assigned the resource (using Algorithm 1).

**Theorem 5** *For any  $m \geq 1$  and constant  $\lambda \in (0, \frac{1}{4})$ , MaxT admits a polynomial time  $\Omega(1)$ -approximation on any instance  $\mathcal{I} = (J, M, \mathcal{W})$  with slackness parameter  $\lambda$ .*

**Proof** Given such an instance  $\mathcal{I}$ , consider the resulting laminar instance. As before, we handle separately two subsets of jobs.

**Subset 1:** For jobs  $j_\ell \in J_\ell$  satisfying  $s_{j_\ell} \leq \alpha_\ell$ , where  $\alpha_\ell$  is defined in (4), apply Algorithm 1 with  $\omega$  value as in (5). Then, the approximation ratio is  $\omega = (1 - \lambda_\ell)^2 \geq (1 - 4\lambda)^2$ .

**Subset 2** For jobs  $j_\ell$  where  $s_{j_\ell} > \alpha_\ell$ , use Lemma 2 to obtain  $\Omega(\frac{1}{\log(1/\alpha_\ell)})$ -approximation.

Taking the best among the solutions for the two subsets of jobs, we obtain an  $\Omega(1)$ -approximation.  $\square$

Finally, consider a general instance of MaxTBy selecting  $\delta = \frac{1}{n}$ , we can apply Lemma 2 to obtain an  $\Omega(\frac{1}{\log n})$ -approximate solution,  $S_1$  for the jobs  $j \in J$  of heights  $s_j \geq \frac{1}{n}$ . Let  $S_2$  be a solution consisting of all jobs  $j$  for which  $s_j < \frac{1}{n}$ . Note that this solution is feasible since  $\sum_{j \in S_2} s_j < 1$ . Selecting the highest profit solution between  $S_1$  and  $S_2$ , we have:

**Corollary 3** *There is a polynomial time  $\Omega(\frac{1}{\log n})$ -approximation algorithm for MaxT instances with arbitrary slackness parameter  $\lambda \in (0, 1]$ .*

### 4 Maximizing utilization

In this section, we obtain an  $\Omega(1)$ -approximation for MaxT instances where the weight of a job is equal to its area. In other words, the goal is to maximize resource utilization. Our result builds on an algorithm of Chen et al. (2002).

**Theorem 6** *There is a polynomial time  $\Omega(1)$ -approximation for any instance of MaxT where  $w_j = a_j$  for all  $j \in J$ .*

**Proof** As before, we represent a time slot  $t$  on a host  $i$  by a bin  $(i, t)$ . We partition the jobs according to their height. Fix some constant  $\delta \in (0, 1)$ . Below, we show how to find a constant approximation for the jobs  $j \in J$  of heights  $s_j \leq \delta$ . For the jobs  $j \in J$  of heights  $s_j > \delta$ , we can obtain a constant approximation using Lemma 2. These two algorithms imply a polynomial time  $\Omega(1)$ -approximation for any instance.

For the rest of this section we assume that the height of every job  $j \in J$  is upper bounded by  $\delta$ . Fix some  $\lambda < \frac{1}{4}$ . Partition the jobs according to their lengths. A job  $j$  is long if  $p_j > \lambda|\chi_j|$ ; otherwise, job  $j$  is short. For a given optimal solution of the problem, let  $OPT_\ell$  and  $OPT_s$  be the contributions of long jobs and short jobs, respectively. We handle the long and short jobs separately. Note that since short jobs satisfy the requirements of Theorem 5, we can obtain a constant approximation with respect to  $OPT_s$ .

We now handle the long jobs. For this part, we adapt an algorithm due to Chen et al. (2002). Let  $L$  be the set of long jobs. Consider the following algorithm.

**Step 1** Sort the jobs in non-increasing order of their time window sizes  $|\chi_j|$ .

**Step 2** For each job  $j$  in the sorted order, if there are  $p_j$  time slots that have at least one bin that is less than  $1 - s_j$  full, schedule  $j$  in these  $p_j$  time slots; otherwise, discard it.

Let  $A$  be the set of jobs scheduled by this algorithm. We now analyze the performance of the algorithm. For each job  $j$ , we define an augmented job  $j'$  as follows:

$$p_{j'} = 3|\chi_j|, \text{ and } s_{j'} = \frac{s_j}{1 - \delta}$$

$$\chi_{j'} = [r_j - |\chi_j|, d_j + |\chi_j|] = [2r_j - d_j, 2d_j - r_j]$$

Let  $A'$  denote the set of augmented jobs for  $A$ . We note that there may be no feasible schedule of the jobs in  $A'$ . Define the weight of  $A'$  as

$$W(A') = \sum_{t \in T} \sum_{j' \in A': t \in \chi_{j'}} s_{j'}$$

It follows that the throughput of the algorithm, denoted  $W(A)$ , satisfies

$$W(A) = \sum_{j \in A} p_j s_j \geq \frac{(1 - \delta)\lambda}{3} W(A')$$

To complete the proof, we simply show that  $W(A') \geq W(OPT_\ell)$ . To this end, it suffices to show that for every  $t \in T$

$$\sum_{j' \in A': t \in \chi_{j'}} s_{j'} \geq \sum_{\omega \in OPT_\ell: t \in \chi_\omega} s_\omega$$

Fix any  $t \in T$ . Two cases arise:

**Case 1** All jobs  $\omega \in OPT_\ell : t \in \chi_\omega$  are scheduled by our algorithm. In this case the claim follows trivially.

**Case 2** At least one long job  $\omega \in OPT_\ell : t \in \chi_\omega$  is rejected by our algorithm. We show that in this case  $\sum_{j' \in A': t \in \chi_{j'}} s_{j'} \geq m$ . The proof follows since  $\sum_{\omega \in OPT_\ell: t \in \chi_\omega} s_\omega \leq m$ .

Since job  $\omega$  was rejected  $\exists t' \in \chi_\omega$  such that each bin  $(i, t')$  is at least  $(1 - \delta)$  full when job  $\omega$  is considered by the algorithm. Let  $A_\omega$  be the set of jobs already scheduled by the algorithm in time slot  $t'$  before  $\omega$  is rejected, and let  $A'_\omega$  be the respective set of augmented jobs. We claim that for every job  $j' \in A'_\omega$ , we have  $t \in \chi_{j'}$ . To see this, note that for any  $j \in A_\omega$ , we have  $|\chi_j| \geq |\chi_\omega|$  (because the jobs are chosen in non-increasing order of their time window sizes). Further,  $\chi_j \cap \chi_\omega$  contains at least  $t'$  and hence  $\chi_j$  and  $\chi_\omega$  are intersecting. Therefore, the time window of the augmented job  $j' \in A'_\omega$  must completely contain  $\chi_\omega$ , and so  $t \in \chi_{j'}$ . Thus, we have

$$\sum_{j' \in A': t \in \chi_{j'}} s_{j'} \geq \sum_{j' \in A'_\omega} s_{j'} \geq \frac{1}{1 - \delta} \sum_{j \in A_\omega} s_j \geq \frac{m(1 - \delta)}{1 - \delta} = m.$$

$\square$

## 5 Resource minimization

In this section, we consider the MinR problem with  $d$  resources, where  $d \geq 1$  is some constant. We show that the problem admits an  $O(\log d)$ -approximation algorithm under some mild assumptions on the slack and minimum window size. Further, we show that the latter assumption can be removed with a slight degradation in the approximation ratio.

Our approach builds on a formulation of the problem as a configuration LP and involves two main phases: a *maximization phase* and *residual phase*.

### 5.1 Configuration LP

We start by describing the configuration linear program that is at the heart of our algorithm. Let  $J_t \subseteq J$  denote the set of all jobs  $j$  such that  $t \in \chi_j$ , i.e.,  $j$  can be allocated resources at time slot  $t$ . For any  $t \in [T]$  and  $S \subseteq J_t$ ,  $C = (S, t)$  is a valid configuration on a single host if  $\sum_{j \in S} \bar{s}_j \leq \bar{I}^d$ , i.e., the jobs in  $S$  can be feasibly allocated their resource requirements on a single host at time slot  $t$ . Denote the set of all valid configurations at time  $t$  by  $C_t$ , and by  $C^j$  the set of all valid configurations  $(S, t)$ , such that  $S$  contains job  $j$ . Denote by  $x_C$  the indicator variable for choosing configuration  $C$ , and by  $m$  the number of hosts needed to schedule all jobs. The fractional relaxation of the Integer Program formulation of our problem is given below.

**Primal :** Minimize  $m$   
 Subject to:  $m - \sum_{C \in C_t} x_C \geq 0, \quad \forall t \in [T]$   
 $-\sum_{C \in C^j \cap C_t} x_C \geq -1, \quad \forall j \in J, t \in [T]$   
 $\sum_{C \in C^j} x_C \geq p_j, \quad \forall j \in J$   
 $x_C \geq 0, \quad \forall C$

The first constraint ensures that we do not pick more than  $m$  configurations for each time slot  $t \in [T]$ . The second constraint guarantees that at most one configuration is chosen for each job  $j$  at a given time  $t$ . Finally, the last constraint guarantees that each job  $j$  is allocated the resource for  $p_j$  time slots, i.e., job  $j$  is completed.

Given that the LP has an exponentially number of variables, we consider solving the dual.

**Dual :** Maximize  $\sum_{j \in J} (p_j \alpha_j - \sum_{t \in [T]} \beta_{j,t})$   
 Subject to:  $\sum_{j \in S} (\alpha_j - \beta_{j,t}) - \gamma_t \leq 0, \quad \forall C = (S, t)$   
 $\sum_{t \in [T]} \gamma_t \leq 1,$   
 $\alpha_j, \beta_{j,t}, \gamma_t \geq 0$

The proof of the following theorem is similar to a result due to Fleischer et al. (2011), with some differences due to the “negative” terms in the objective of the dual program.

**Theorem 7** For any  $\epsilon > 0$ , there is a polynomial time algorithm that yields a  $(1 + \epsilon)$ -approximate solution for the configuration LP.

**Proof** We first describe a separation oracle for the dual program. Given a vector  $(\gamma_t : t \in [T], \beta_{j,t} : j \in J, t \in [T], \alpha_j : j \in J)$ , the oracle should either report that the solution is feasible or find a violating constraint. Clearly, the non-trivial task is to check the exponentially many constraints corresponding to the configurations. To this end, we compute a configuration  $C = (S, t)$ , for each  $t$ , that maximizes the value  $\sum_{j \in S} (\alpha_j - \beta_{j,t})$ . Subsequently, we can compare this value against the fixed value  $\gamma_t$ . Further, observe that such a subset  $S$  (for a given  $t$ ) can be approximately found by solving the following instance of the *multi-dimensional knapsack* problem. Indeed, we have an item for each job  $j \in J_t$  with size  $s_j$  and profit  $\alpha_j - \beta_{j,t}$ . The goal is to find a subset of items of maximum total profit that can fit into the  $d$ -dimensional bin  $1^d$ . There is a well known PTAS for the multi-dimensional knapsack problem, for any fixed  $d > 1$  (Frieze & Clarke, 1984). Let  $\epsilon' > 0$  be the error parameter for the PTAS.

We now run the ellipsoid algorithm on the dual program. We perform a binary search over the possible values of  $z^* \leq \sum_{j \in J} (\alpha_j p_j - \sum_{t \in [T]} \beta_{j,t})$ . For  $z^*$ , the ellipsoid algorithm with the given separation oracle reports success.

Suppose that the ellipsoid algorithm reports failure at the value  $z^* + \delta$ , where  $\delta$  is the accuracy parameter of the binary search, which can be made as small as desired. Clearly, the optimal solution value must be lower than  $z^* + \delta$ . On the other hand, since we are using a  $(1 + \epsilon')$ -approximation oracle, we have that if a solution  $(\alpha_j, \beta_{j,t}, \gamma_t)$  is reported to be feasible, then we must have that  $(\alpha_j / (1 + \epsilon'), \beta_{j,t} / (1 + \epsilon'), \gamma_t)$  is feasible for the original dual program. Hence, the optimal solution lies in the range  $[z^* / (1 + \epsilon'), z^* + \delta]$ .

Further, we look at the constraints checked by the ellipsoid algorithm when the value is set to be  $z^* + \delta$ . There are polynomial number of such constraints before the algorithm reports a failure. We consider the dual of this restricted LP that is equivalent to a restricted original configuration LP obtained by setting to zero the variables corresponding to the constraints not considered by the ellipsoid algorithm. As noted by Fleischer et al. (2011), the cost of this LP is at most  $z^* + \delta$  by LP duality. Thus, for appropriate selection of  $\epsilon' \in (0, \epsilon)$ , we obtain a  $(1 + \epsilon)$ -approximate solution for the configuration LP.  $\square$

### 5.2 The algorithm

Let  $m^*$  denote the optimal value of the configuration LP, and let  $m \leq (1 + \epsilon)m^*$  be the objective value of the approximate solution of the configuration LP, rounded up to the nearest integer. The detailed description of the algorithm is

**Input:** MinR instance  $(J, \mathcal{W})$ , where  $J$  is the set of jobs and  $\mathcal{W}$  is the set of time-windows

**Output:** An assignment of all jobs in  $J$  to a (approximate) minimum number of hosts

- 1: Solve the configuration LP approximately using Theorem 7 to obtain a (fractional) solution  $\{\hat{x}_C : \forall C\}$ . Let  $m$  denote the objective value of the solution (assume  $m$  is rounded up to the nearest integer).
- 2: Let  $c > 1$  be a constant (to be determined). For each  $t \in [T]$  repeat: for  $m_1 = cm \log d$  iterations, choose a configuration  $C \in \mathcal{C}_t$  with probability  $\hat{x}_C/m$ . If two configurations  $C_1 = (S_1, t)$  and  $C_2 = (S_2, t)$  where  $S_1 \cap S_2 \neq \emptyset$  are chosen in this process, replace  $C_2$  with  $C'_2 = (S_2 \setminus S_1, t)$ . Continue with this replacement process until all configuration sets corresponding to a given time-slot  $t$  are disjoint.
- 3: If a configuration  $C = (S, t)$  is chosen in Step 2 in iteration  $i$  for time  $t$ , assign the jobs in  $S$  to host  $i$  at time  $t$ .
- 4: Let  $J_{res}$  be the set of residual jobs constructed from  $J$  as follows: for each  $j \in J$ , let  $n_j$  denote the number of configurations containing job  $j$  that are chosen in Step 2. Now, define a residual job  $j'$  with attributes  $(p_{j'} = \max(p_j - n_j, 0), s_{j'} = \|\bar{s}_j\|_\infty, \chi_{j'} = \chi_j)$ . For each job  $j'$  associate a set,  $forb(j')$ , of forbidden time slots in which a configuration containing job  $j$  is already chosen.
- 5: Use the construction given in Sect. 3.2 to transform  $\mathcal{W}$  into a laminar family of intervals  $\mathcal{L}$  and obtain the corresponding mapping  $\mathfrak{L} : \mathcal{W} \rightarrow \mathcal{L}$ . Consider a modified instance  $(J'_{res} = \{j' : j \in J_{res}\}, \mathcal{L})$  such that  $\chi_{j'} = \mathfrak{L}(\chi_j)$ . For each  $j' \in J'_{res}$  let  $forb(j')$  be the forbidden time slots in  $\chi_{j'}$ .
- 6: Use Lemma 6 to compute a feasible schedule for the residual jobs  $J'_{res}$ . Let  $m_2$  be the number of additional hosts needed in this step.
- 7: Output the resulting assignment of the jobs in  $J$  to  $m_1 + m_2$  hosts.

**Algorithm 4:** Algorithm for the MinR problem

given in Algorithm 4. We use the following two stage process. In the first stage (Steps 2 and 3 of Algorithm 4), we choose  $O(m \log d)$  configurations  $C$  (with probabilities proportional to their  $x_C$  values in the LP solution) for each time slot. Indeed, this random selection may lead to *partial* execution of some jobs  $j$ , which are allocated the resources for less than  $p_j$  time slots. The second stage (Steps 4–6 of the algorithm) amends this, by considering the “residual” job parts and assigning them to a set of  $m$  new hosts.

Our key technical result (in Lemma 5) is that, with high probability, the total volume of the residual jobs to be scheduled in any time window  $\chi$  is sufficiently small. Some additional challenges arise due to the fact that the time slots used to schedule a job in the first and second stages must be disjoint. Thus, the time slots used for job  $j$  in the first stage become “forbidden” for  $j$  in the second stage. We associate with each job  $j$  a subset of “forbidden” time slots, denoted  $forb(j) \subseteq \chi_j$ . Any feasible solution for the residual jobs must ensure that job  $j$  is not scheduled at time  $t \in forb(j)$ . To resolve this issue we need to refine Algorithm 2 (in Lemma 6).

**5.3 Analysis**

Towards analyzing our algorithm, we prove several technical lemmas.

**Definition 1** A real valued function  $f : (\mathcal{M}_1 \times \mathcal{M}_2 \times \dots \times \mathcal{M}_\ell) \rightarrow \mathbb{R}$  satisfies the *bounded differences property* if there are  $\ell$  constants  $c_1, \dots, c_\ell \in \mathbb{R}$  such that  $\forall \bar{\mu}, \bar{\mu}' \in \mathcal{M}_1 \times \dots \times \mathcal{M}_\ell$ , if  $\bar{\mu}$  and  $\bar{\mu}'$  differ in exactly the  $i$ th coordinate then  $|f(\bar{\mu}) - f(\bar{\mu}')| \leq c_i$ .

The following useful result is due to McDiarmid (1989).

**Lemma 4** (McDiarmid) *Let  $Y = (Y_1, Y_2, \dots, Y_\ell)$  be a family of independent random variables, such that the sample space (domain) of  $Y_j$  is  $\mathcal{M}_j$ , and let  $f : (\mathcal{M}_1 \times \dots \times \mathcal{M}_\ell) \rightarrow \mathbb{R}$  be a function whose domain is the  $\ell$ -dimensional sample space that satisfies the bounded differences property. Then,  $Pr[f(\bar{\mu}) - \mathbb{E}(f(\bar{\mu})) \geq \psi] \leq \exp(-2\psi^2 / \sum_{i=1}^\ell c_i^2)$ .*

The next lemma gives the conditions implying that the total volume of the residual jobs is *small*. This is essential for obtaining our performance bounds.

**Lemma 5** *Let  $\chi$  be any sub-interval of  $[T]$ . For any  $\epsilon \in (0, 1)$ ,  $\omega \in (0, 1)$ , and a sufficiently large value of  $\theta$  that depends on  $\epsilon$  and  $\omega$ , the following holds with probability at least  $1 - \epsilon$*

$$\sum_{j \in J_{res}: \chi_j \subseteq \chi} p'_j \|\bar{s}_j\|_\infty \leq \omega m |\chi|, \tag{6}$$

assuming interval  $\chi$  satisfies

$$|\chi| \geq \frac{1}{m} \theta d^2 \log d \log(\epsilon^{-\frac{1}{2}} T), \tag{7}$$

where  $d > 1$  is the dimension of the vectors  $\bar{s}_j$ ,  $j \in J$ . Further, if  $T = O(d^\delta)$ , for some constant  $\delta \geq 0$ , then the restriction on the length of  $\chi$  in (7) can be dropped and (6) holds for any interval  $\chi$ .<sup>7</sup>

**Proof** Consider an interval  $\chi$  that satisfies (7) and a job  $j$  such that  $\chi_j \subseteq \chi$ . Define  $X_{jt}$  as the total (fractional) value of configurations corresponding to job  $j$  and time  $t$ , i.e.,  $X_{jt} = \sum_{C \in \mathcal{C}^j \cap \mathcal{C}_t} \hat{x}_C$ . Then,

$$\sum_{t \in \chi_j} X_{jt} = \sum_{C \in \mathcal{C}^j} \hat{x}_C \geq p_j. \tag{8}$$

Further, by the LP constraints, we have that  $X_{jt} \leq 1$ .

For the analysis, we partition  $\chi$  into  $p_j$  regions  $R_1^j, R_2^j, \dots, R_{p_j}^j$ , such that

$$\sum_{t \in R_k^j} X_{jt} \geq 1/2. \tag{9}$$

<sup>7</sup> As we show in the proof, we implicitly use here the constant  $c$  in Step 2 of Algorithm 4, which depends on  $\omega$ . For the case where  $T = O(d^\delta)$ ,  $c$  depends on  $\epsilon$ ,  $\omega$ , and  $\delta$ .



The regions are not necessarily formed of consecutive time slots in  $\chi$ . One way of doing this is as follows. Start with singleton regions  $\{t\}$  for which  $X_{jt} \geq 1/2$ . Let  $q_j$  denote the number of regions generated. If  $q_j \geq p_j$  then we are done; otherwise, until  $p_j$  regions are obtained, we do the following, for  $k \in [q_j + 1, p_j]$ : starting with  $R_k^j = \emptyset$  and while  $\sum_{t \in R_k^j} X_{jt} < 1/2$ , add a time slot  $t'$  that has not yet been assigned to any region. Since any such  $t'$  has not been chosen as a singleton region, we have  $X_{jt'} < 1/2$ ; therefore,  $\sum_{t \in R_k^j} X_{jt} \leq 1$ . From Equation (8) we have that at least  $p_j$  regions are generated in this process.

Let  $\mathbf{R}_k^j$  denote the event that job  $j$  is not scheduled in a given region  $R_k^j$  in the first stage of the algorithm. We now compute the probability  $\mathbb{P}(\mathbf{R}_k^j)$ . Let  $Pr(j, t)$  be the probability that  $j$  is not assigned to any host at time  $t$ . Then,

$$Pr(j, t) = \left(1 - \frac{X_{jt}}{m}\right)^{cm \log d} \leq \exp(-cX_{jt} \log d), \tag{10}$$

where  $c > 1$  is the constant in Step 2 of Algorithm 4. Hence, the probability that job  $j$  is not scheduled in region  $R_k^j$  satisfies

$$\begin{aligned} \mathbb{P}(\mathbf{R}_k^j) &\leq \prod_{t \in R_k^j} \exp(-cX_{jt} \log d) \\ &= \exp(-c \log d \sum_{t \in R_k^j} X_{jt}) \end{aligned} \tag{11}$$

$$\leq \exp\left(-\frac{c}{2} \log d\right) = d^{-\frac{c}{2}} \tag{12}$$

The last inequality follows from (9).

Recall that  $p'_j$  (in Algorithm 4) is the processing time required for the residual job  $j$ . The value of  $p'_j$  is upper bounded by the number of regions in which job  $j$  has not been scheduled. Therefore, we have

$$\mathbb{E}[p'_j] \leq \sum_{k \in [p_j]} \mathbb{P}(\mathbf{R}_k^j) \leq p_j d^{-\frac{c}{2}}. \tag{13}$$

By the LP constraints, we have that  $\sum_{j: \chi_j \subseteq \chi} p_j \bar{s}_j \leq m|\chi| \bar{1}^d$ . We note that each bin  $(i, t)$  for  $t \in \chi$  and  $i \geq 1$  is a cube  $\bar{1}^d$  in which we “pack” a subset of jobs  $j$  satisfying  $\chi_j \subseteq \chi$ . For a job  $j$  packed in  $(i, t)$ , let  $1 \leq \eta \leq d$  be the resource for which  $\bar{s}_{j,\eta} = \|\bar{s}_j\|_\infty$ . Now, replace job  $j$  by a job  $j'$  with a single resource requirement, given by  $\|\bar{s}_j\|_\infty$ . Then we can replace the bin  $(i, t)$  by  $d$  unit size bins and assign  $j'$  into bin  $\eta$ . Since  $1 \leq \eta \leq d$ , and due to the feasibility of the packing in  $(i, t)$ , after this transformation, all of the jobs packed in  $(i, t)$  can be feasibly assigned into the  $d$  unit size bins. It follows that  $\sum_{j: \chi_j \subseteq \chi} p_j \|\bar{s}_j\|_\infty \leq m|\chi|d$ . Hence, using (13) we obtain the following upper bound on

the expected total volume of residual jobs in  $\chi$ :

$$\mathbb{E} \left[ \sum_{j: \chi_j \subseteq \chi} p'_j \|\bar{s}_j\|_\infty \right] \leq m|\chi|d^{1-\frac{c}{2}} \tag{14}$$

To complete the proof, we need to show that, with high probability, the total volume of residual jobs in any sub-interval  $\chi$  of  $[T]$  is small. To this end, we apply Lemma 4 as follows. For each  $t \in \chi$  and iteration  $i \in [cm \log d]$ , there is an associated random variable  $Y_{t,i}$  indicating the configuration  $C$  chosen in iteration  $i$  at time  $t$ . Note that since Lemma 4 requires independence between events, the configurations considered are the ones prior to the modifications applied in Step 2. This is valid as we define the function  $f(\cdot)$  in Lemma 4 to take into account the modifications applied in Step 2. This function is defined as the quantity  $A_{res}[\chi] = \sum_{j: \chi_j \subseteq \chi} p'_j \|\bar{s}_j\|_\infty$  which is a function of these  $cm|\chi| \log d$  independent random variables  $\{Y_{t,i}\}$ . If one of these variables is altered, it might affect the choice of at most one configuration, namely, a configuration  $C'$  selected in iteration  $i$  at time  $t$  is replaced by another configuration  $C''$ . Suppose  $A'_{res}(\chi)$  and  $A''_{res}(\chi)$  are the two corresponding realizations of the random variable  $A_{res}(\chi)$ . We now bound the quantity  $|A'_{res}(\chi) - A''_{res}(\chi)|$ : the worst scenario is clearly when none of the jobs in configuration  $C'$  is contained in any other configuration chosen at time  $t$ , whereas every job in  $C''$  is contained in some other configuration chosen at time  $t$ ; or vice versa. Thus, we have  $|A'_{res}(\chi) - A''_{res}(\chi)| \leq \max(\sum_{j \in C'} \|\bar{s}_j\|_\infty, \sum_{j \in C''} \|\bar{s}_j\|_\infty) \leq d$ . Applying Lemma 4, we have, for any  $\omega' \geq 0$

$$\begin{aligned} Pr[A_{res}(\chi) - \mathbb{E}(A_{res}(\chi)) \geq \omega' m|\chi|] &\leq \exp\left(\frac{(-2\omega'^2)m^2|\chi|^2}{cm \log d|\chi|d^2}\right) \\ &= \exp\left(-\frac{2\omega'^2}{c} \cdot \frac{m|\chi|}{d^2 \log d}\right) \end{aligned}$$

Let  $c > 2$  be a sufficiently large constant such that  $\omega > d^{1-\frac{c}{2}}$ . We set  $\omega' = \omega - d^{1-\frac{c}{2}}$ , and  $\theta = c\omega'^{-2}$ . Then, we have

$$\begin{aligned} Pr[A_{res}(\chi) \geq \omega m|\chi|] &= Pr[A_{res}(\chi) \geq (\omega' + d^{1-\frac{c}{2}})m|\chi|] \\ &\leq Pr[A_{res}(\chi) - \mathbb{E}(A_{res}(\chi)) \geq \omega' m|\chi|] \\ &\leq \exp\left(-\frac{2\omega'^2}{c} \cdot \frac{m|\chi|}{d^2 \log d}\right) = \exp\left(\frac{-2}{\theta} \cdot \frac{m|\chi|}{d^2 \log d}\right) \\ &\leq \exp(-2 \log(\epsilon^{-\frac{1}{2}} T)) \leq \epsilon T^{-2} \end{aligned}$$

The first inequality follows from (14) and the third from (7). Since the total number of distinct intervals possible in  $[T]$  is at most  $T^2$ , by applying the union bound, the probability that

some interval  $\chi$  that satisfies (7) fails to satisfy (6) is at most  $\epsilon$ .

Now, consider the case where  $T = d^\delta$ , for some constant  $\delta \geq 0$  (independent of  $d$ ). We have

$$\begin{aligned} Pr[A_{res}(\chi) \geq \omega m|\chi|] &\leq Pr[A_{res}(\chi) \geq \frac{\mathbb{E}(A_{res}(\chi))\omega}{d^{1-\frac{\epsilon}{2}}}] \\ &\leq \frac{d^{1-\frac{\epsilon}{2}}}{\omega}. \end{aligned}$$

The first inequality follows from (14) and the second from Markov’s inequality. As before, we observe that total number of distinct intervals possible in  $[T]$  is at most  $T^2 \leq d^{2\delta}$ . Now, we choose  $\epsilon' > 0$  to be a large enough constant satisfying  $\omega^{-1}d^{-\frac{\epsilon'}{2}} \leq \epsilon$ , and  $c \geq 2(2\delta + 1) + \epsilon'$ . Hence,

$$Pr[A_{res}(\chi) \geq \omega m|\chi|] \leq \frac{d^{1-(2\delta+1)-\epsilon'/2}}{\omega} \leq \epsilon T^{-2}.$$

Again, since the total number of distinct intervals possible in  $[T]$  is at most  $T^2$ , by applying the union bound, the probability that some interval  $\chi$  fails to satisfy (6) is at most  $\epsilon$ .  $\square$

The next lemma gives the conditions that guarantee feasible schedule for the residual jobs.

**Lemma 6** *Let  $(J, \mathcal{L})$  be a single resource laminar instance of MinR such that for any job  $j$ ,  $p_j + |\text{forb}(j)| \leq \lambda|\chi_j|$ , for some  $\lambda \in (0, 1)$ , and for any  $\chi \in \mathcal{L}$ :  $\sum_{\{j \in J: \chi_j \subseteq \chi\}} p_j s_j \leq \alpha m|\chi|$ , for some  $\alpha \in (0, 1)$ . Then, if  $\lambda + 2\alpha \leq 1$ , all jobs in  $J$  can be assigned to  $m + 1$  hosts.*

**Proof** We use an algorithm similar to Algorithm 2. The only difference is in the definition of  $\text{avail}(j)$ . In our case  $\text{avail}(j) \subseteq \chi_j \setminus \text{forb}(j)$  since job  $j$  cannot be assigned to bins in any of the time slots in  $\text{forb}(j)$ . Recall that Algorithm 2 works in a bottom-up fashion and a job  $j$  is processed after all intervals  $\chi' \subset \chi_j$  have been marked *done*. While processing job  $j$ , a bin is labelled with one of three possible colors: *white*, *gray* or *black*. Initially, all bins are colored *white*.

Similar to the proof of Theorem 2, we prove that the algorithm never reports **fail** while processing job  $j$ . Assume towards a contradiction that the algorithm reports **fail** while assigning job  $j$ . We say that time slot  $t \in \chi_j$  is *unavailable* if either  $t \in \text{forb}(j)$  or  $j$  has been assigned to some bin  $(i, t)$  in this time slot. Otherwise, time slot  $t \in \chi_j$  is *available*. Since  $p_j + |\text{forb}(j)| \leq \lambda|\chi_j|$  the number of unavailable time slots while processing job  $j$  is at most  $\lambda|\chi_j|$ . We say that time slot  $t \in \chi_j$  is *bad* if it contains a gray bin. Similar to the proof of Theorem 2 (although the definition of a bad time slot here is slightly different), it can be proven that as long as no job  $j^+$  such that  $\chi_j \subset \chi_{j^+}$  has been allocated the resource: the number of bad time slots while processing job  $j$  is at most

$\lambda|\chi_j|$ . The only difference in the proof is in case bin  $(i, t)$  becomes gray after the assignment. In this case, since there may be bad time slots in  $\text{forb}(j)$ , we need to use the inequality  $p_j + |\text{forb}(j)| \leq \lambda|\chi_j|$  (rather than  $p_j \leq \lambda|\chi_j|$ ).

Since we pair the black bins  $(i, t) \leftrightarrow (i', t')$  only if  $\sum_{j \in \mathcal{S}(i,t)} s_j + \sum_{j' \in \mathcal{S}(i',t')} s_{j'} > 1$ , the total number of black bins  $< 2\alpha m|\chi_j|$ . Hence, the total number of bins that are black or in unavailable time slots is  $< (\lambda + 2\alpha)m|\chi_j|$ . Since a time slot  $t$  may contain at most one gray bin, the number of gray bins in available time slots is at most  $\lambda|\chi_j|$ . Thus, if  $\lambda + 2\alpha \leq 1$ , we have  $(\lambda + 2\alpha)m|\chi_j| + \lambda|\chi_j| < (m + 1)|\chi_j|$ , and there should be at least one available time slot  $t^*$  that contains a white bin  $(i^*, t^*)$ . But in this case, we could have assigned  $j$  to host  $i^*$  at time  $t^*$ , which is a contradiction to the assumption that the algorithm reports a **fail**.  $\square$

The above lemmas lead to an  $O(\log d)$  performance guarantee for instances with large time windows, as formalized in the next result.

**Theorem 8** *Let  $(J, \mathcal{W})$  be an instance of MinR with slackness parameter  $\lambda \in (0, \frac{1}{4})$ . Fix an  $\epsilon \in (0, 1)$ . If  $|\chi_j| \geq \frac{1}{m}\theta d^2 \log d \log(T\epsilon^{-\frac{1}{2}}) \forall j \in J$ , for sufficiently large constant  $\theta = \theta(\epsilon)$ , then Algorithm 4 yields an  $O_\epsilon(\log d)$  approximation ratio<sup>8</sup> with probability at least  $1 - \epsilon$ .*

**Proof** The optimal objective value of the configuration LP, denoted  $m^*$  is a lower bound on the number of hosts required for the instance  $(J, \mathcal{W})$ . By Theorem 7,  $m \leq (1 + \epsilon)m^*$ , where  $m$  is the (rounded) objective value of the approximate solution for the configuration LP. Now, assuming Algorithm 4 is correct, the number of hosts used is at most  $m(c \log d) + m$ , implying an  $O(\log d)$  approximation ratio. Below, we prove correctness of Algorithm 4, i.e., we show that it feasibly schedules all the jobs in  $J$ .

It suffices to show that the residual set of jobs  $J_{res}$  can be successfully scheduled. Recall that Algorithm 4 uses the construction given in Sect. 3.2 to transform  $\mathcal{W}$  into a laminar set of intervals  $\mathcal{L}$  and to obtain a mapping  $\mathcal{L} : \mathcal{W} \rightarrow \mathcal{L}$ . Then, the algorithm computes a schedule of the residual set of jobs  $J_{res}$  by solving the laminar instance  $(J'_{res} = \{j' : j \in J_{res}\}, \mathcal{L})$ , where  $\chi_{j'} = \mathcal{L}(\chi_j)$ . Observe that for any job  $j \in J_{res}$  we have that  $p'_j + |\text{forb}(j)| = p_j \leq \lambda|\chi_j| \leq 4\lambda|\chi_{j'}|$ , where the last inequality follows from Lemma 1. By the same lemma, we also get that for any job  $j \in J_{res}$ ,  $|\chi_{j'}| \geq \frac{1}{4}|\chi_j|$ . Now, let  $\tilde{\chi}_{j'}$  be the union of all the time windows mapped by  $\mathcal{L}$  to time windows in  $\chi_{j'}$ . Also, by Lemma 1  $|\tilde{\chi}_{j'}| \leq 4|\chi_{j'}|$ . Clearly,  $|\tilde{\chi}_{j'}| \geq |\chi_{j'}| \geq \frac{1}{4m}\theta d^2 \log d \log(T\epsilon^{-\frac{1}{2}})$ . Let  $\omega \in (0, 1)$ . Next, we apply Lemma 5 to the intervals  $\tilde{\chi}_{j'}$  for  $j \in J_{res}$  (with appropriate value of  $\theta$ ), to get that that  $\sum_{j \in J_{res}: \chi_{j'} \subseteq \chi} p'_j \|\tilde{s}_j\|_\infty \leq 4\omega m|\chi|$  with probability at least

<sup>8</sup> The constant in the Big- $O_\epsilon$  depends on  $\epsilon$ .

$1 - \epsilon$ . We now apply Lemma 6, by setting  $\alpha \leftarrow 4\omega$  and  $\hat{\lambda} \leftarrow 4\lambda$ . We note that there is a feasible schedule of jobs  $j'$  on  $m + 1$  hosts if  $4\lambda + 8\omega = 1$ . Indeed, for any  $\lambda < \frac{1}{4}$ , there is a positive constant  $\omega$  that satisfies this equation. Finally, it is easy to show that transforming the instance back to  $d$  dimensions (by replacing the requirement of  $\|\bar{s}_j\|_\infty$  by  $\bar{s}_j$ ) the schedule remains feasible. This completes the proof.  $\square$

We conclude our analysis with the following result.

**Theorem 9** *Let  $(J, \mathcal{W})$  be an instance of MinR with slackness parameter  $\lambda \in (0, \frac{1}{4})$ . Fix an  $\epsilon \in (0, 1)$ . There is a polynomial time algorithm that yields an  $O_\epsilon(\log d \log^* T)$  approximation ratio with probability at least  $1 - \epsilon$ .*

**Proof** The key idea is the following. Starting with the maximum schedule length  $T$ , we recursively define  $\kappa + 1$  ranges for the time-window sizes in the original instance. We then partition the set of jobs  $J$  to  $\kappa + 1$  subsets, each containing jobs with time windows within the corresponding range, where  $\kappa = O(\log^* T)$ . The crux of this partition is that the resulting  $\kappa + 1$  instances of our problem satisfy the conditions of Theorem 8. In particular, all jobs have ‘large’ windows. Thus, we can obtain for each instance a  $\log d$ -approximate solution. Formally, let  $\theta$  be the constant in Theorem 8. Set  $\gamma = \theta d^2 \log d$  (we assume that the log is base 2). Define the function  $\psi : \mathbf{N} \rightarrow \mathbf{R}^+$  as follows:

$$\psi(i) = \begin{cases} 4\lceil \gamma^2 \rceil & \text{if } i = 1 \\ \min\{T, \sqrt{\epsilon} (2^{1/(2\gamma)})^{\psi(i-1)}\} & \text{if } i > 1 \end{cases}$$

It is easy to verify that, for  $i \geq 1$ , we have  $\psi(i) \leq \psi(i + 1)$ , and  $\psi(i) \geq 2\gamma \log_2 (\psi(i + 1)\epsilon^{-\frac{1}{2}})$ . Also, let  $\kappa$  be the smallest integer for which  $\psi(\kappa) = T$ . Then, we have  $\kappa = O(\log^* T)$ .

We partition the set of intervals  $\mathcal{W}$  into groups based on their length as follows:

$$\begin{aligned} \mathcal{I}_0 &= \{\chi : |\chi| \in [1, \psi(1)]\}, \\ \mathcal{I}_w &= \{\chi : |\chi| \in (\psi(w), \psi(w + 1)]\} \quad \text{for all } w \in [1, \kappa]. \end{aligned}$$

Next, we define  $\kappa + 1$  instances of our problem, where the  $w^{\text{th}}$  instance, for  $w \in [0, \kappa]$  is given by:

$$(J_w = \{j \in J : \chi_j \in \mathcal{I}_w\}, \mathcal{I}_w).$$

Since each of the above instances requires to schedule a subset of jobs in the original instance, they optimally need at most  $m$  hosts to complete all jobs. Consider the  $w^{\text{th}}$  instance, for  $w \geq 0$ . The largest window size here is at most  $\psi(w + 1)$ . We further partition this instance as follows. Let

$$\begin{aligned} \mathcal{I}_{w,Odd}^i &= [(2i + 1)\psi(w + 1) + 1, \\ &\min\{(2i + 3)\psi(w + 1), T\}], \end{aligned}$$

where  $i \geq 0$  and  $(2i + 1)\psi(w + 1) < T$ . Similarly, let

$$\mathcal{I}_{w,Even}^i = [(2i)\psi(w + 1) + 1, \min\{(2i + 2)\psi(w + 1), T\}],$$

where  $i \geq 0$  and  $(2i)\psi(w + 1) < T$ . Now, we define

$$J_{w,Odd}^i = \{j \in J_w : \chi_j \subseteq \mathcal{I}_{w,Odd}^i\}$$

and

$$J_{w,Even}^i = \{j \in J_w : \chi_j \subseteq \mathcal{I}_{w,Even}^i\}.$$

Let  $J_{w,Odd} = \cup_i J_{w,Odd}^i$  and  $J_{w,Even} = \cup_i J_{w,Even}^i$ . Finally, remove each job  $j \in J_{w,Odd} \cap J_{w,Even}$  from  $J_{w,Even}$  and the corresponding  $J_{w,Even}^i$ . Consequently,  $J_{w,Odd} \cap J_{w,Even} = \emptyset$ .

In the analysis we distinguish between the case  $w > 0$  and  $w = 0$ . For any  $w > 0$ , fix an  $i \geq 0$  such that  $(2i + 1)\psi(w + 1) < T$  and consider the instance defined by the jobs in  $J_{w,Odd}^i$ . We claim that this instance can be solved using Theorem 8. Indeed, the total number of time slots in this instance is  $2\psi(w + 1)$ , and the time-window of any job  $j$  in the instance is  $|\chi_j| \geq \psi(w) \geq 2\theta d^2 \log d \log (\psi(w + 1)\epsilon^{-\frac{1}{2}})$ . Thus, the conditions in Theorem 8 are satisfied, and we can obtain a feasible schedule using  $O(m \log d)$  hosts.

For  $w = 0$ , fix an  $i \geq 0$  such that  $(2i + 1)\psi(1) < T$  and consider the instance defined by the jobs in  $J_{0,Odd}^i$ . We claim that this instance can also be solved using Theorem 8, since the total number of time slots in this instance is  $2\psi(1) = O(d^\delta)$ , for some constant  $\delta \geq 0$ , and  $d$  is fixed.

Note that for any  $w$ , the odd instances  $J_{w,Odd}^i, J_{w,Odd}^\ell$  for  $i \neq \ell$  are mutually disjoint (jobs and time-windows). Thus, we can solve them in parallel using the same  $O(m \log d)$  hosts. We can do the same for  $J_{w,Even}$ . Suppose we need  $m_o$  and  $m_e$  hosts to solve  $J_{w,Odd}$  and  $J_{w,Even}$ , respectively. Since no job is shared between  $J_{w,Odd}$  and  $J_{w,Even}$ , we can schedule all the jobs in  $J_w$  using  $m_o + m_e = O(m \log d)$  hosts.

Now, to handle the instances corresponding to all  $w \in \{0, \dots, \kappa\}$ , we note that no job is shared among the instances. Therefore, we can aggregate the hosts to obtain a feasible schedule for all instances using  $O(m\kappa \log d)$  hosts.  $\square$

## 6 Conclusion and open problems

In this work, we address two natural variants of the classic *resource allocation problem* with objectives of *throughput maximization* (MaxT) and *resource minimization* (MinR), respectively. For the MaxT problem, we obtain a constant approximation algorithm assuming there is sufficient slack in completing each job. For the MinR problem, we obtain an algorithm with an approximation ratio that is logarithmic in

the number of resources  $d$ , assuming that none of the time windows is too small. While these assumptions are mild and reasonable in practice, we were unable to remove them. The problem of relaxing these assumptions (or proving its impossibility) remains open.

**Acknowledgements** We thank two anonymous referees for many insightful comments on the paper.

## References

- Adler, M., Gibbons, P. B., & Matias, Y. (2002). Scheduling space-sharing for internet advertising. *Journal of Scheduling*, 5(2), 103–119.
- Ageev, A. A., & Sviridenko, M. (2004). Pipe rounding: A new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 8(3), 307–328.
- Bansal, N., Eliáš, M., & Khan, A. (2016). Improved approximation for vector bin packing. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on discrete algorithms* (pp. 1561–1579).
- Bansal, N., Caprara, A., & Sviridenko, M. (2009). A new approximation method for set covering problems, with applications to multi-dimensional bin packing. *SIAM Journal on Computing*, 39(4), 1256–1278.
- Bansal, N., Friggstad, Z., Khandekar, R., & Salavatipour, M. R. (2014). A logarithmic approximation for unsplittable flow on line graphs. *ACM Transactions on Algorithms*, 10(1), 1.
- Bar-Noy, A., Bar-Yehuda, R., Freund, A., Naor, J., & Schieber, B. (2001). A unified approach to approximating resource allocation and scheduling. *Journal of the ACM*, 48(5), 1069–1090.
- Bar-Noy, A., Guha, S., Naor, J., & Schieber, B. (2001). Approximating the throughput of multiple machines in real-time scheduling. *SIAM Journal on Computing*, 31(2), 331–352.
- Călinescu, G., Chakrabarti, A., Karloff, H. J., & Rabani, Y. (2011). An improved approximation algorithm for resource allocation. *ACM Transactions on Algorithms*, 7(4), 48:1–48:7.
- Chakaravarthy, V. T., Choudhury, A. R., Gupta, S., Roy, S., & Sabharwal, Y. (2014). Improved algorithms for resource allocation under varying capacity. In *European symposium on algorithms* (pp. 222–234).
- Chekuri, C., & Khanna, S. (2004). On multidimensional packing problems. *SIAM Journal on Computing*, 33(4), 837–851.
- Chen, B., Hassin, R., & Tzur, M. (2002). Allocation of bandwidth and storage. *IIE Transactions*, 34(5), 501–507.
- Chuzhoy, J., & Codenotti, P. (2009). Resource minimization job scheduling. In *Approximation, randomization, and combinatorial optimization. Algorithms and techniques* (pp. 70–83).
- Chuzhoy, J., Guha, S., Khanna, S., & Naor, J. (2004). Machine minimization for scheduling jobs with interval constraints. In *Proceedings of the 45th symposium on foundations of computer science (FOCS 2004), 17–19 October 2004, Rome, Italy* (pp. 81–90).
- Dawande, M., Kumar, S., & Sriskandarajah, C. (2003). Performance bounds of algorithms for scheduling advertisements on a web page. *Journal of Scheduling*, 6(4), 373–394.
- Dawande, M., Kumar, S., & Sriskandarajah, C. (2005). Scheduling web advertisements: A note on the MINSPACE problem. *Journal of Scheduling*, 8(1), 97–106.
- Fleischer, L., Goemans, M. X., Mirrokni, V. S., & Sviridenko, M. (2011). Tight approximation algorithms for maximum separable assignment problems. *Mathematical Operations Research*, 36(3), 416–431.
- Fox, K., & Korupolu, M. (2013). Weighted flowtime on capacitated machines. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on discrete algorithms* (pp. 129–143).
- Freund, A., & Naor, J. (2004). Approximating the advertisement placement problem. *Journal of Scheduling*, 7(5), 365–374.
- Frieze, A. M., & Clarke, M. R. B. (1984). Approximation algorithms for the m-dimensional 0–1 knapsack problem: Worst-case and probabilistic analyses. *European Journal of Operational Research*, 15(1), 100–109.
- Jain, N., Menache, I., Naor, J., & Yaniv, J. (2015). Near-optimal scheduling mechanisms for deadline-sensitive jobs in large computing clusters. *ACM Transactions on Parallel Computing*, 2(1), 3.
- Jansen, K., & Porkolab, L. (2002). On preemptive resource constrained scheduling: Polynomial-time approximation schemes. *Integer Programming and Combinatorial Optimization*, 47, 329–349.
- Kalyanasundaram, B., & Pruhs, K. (2001). Eliminating migration in multi-processor scheduling. *Journal of Algorithms*, 38(1), 2–24.
- Kaul, A., Aggarwal, S., Gupta, A., Dayama, N., Krishnamoorthy, M., & Jha, P. C. (2017). Optimal advertising on a two-dimensional web banner. *International Journal of System Assurance Engineering and Management*, 9, 1–6.
- Kumar, S., Dawande, M., & Mookerjee, V. (2007). Optimal scheduling and placement of internet banner advertisements. *IEEE Transactions on Knowledge and Data Engineering*, 19(11), 1571–1584.
- Lawler, E. L. (1990). A dynamic programming algorithm for preemptive scheduling of a single machine to minimize the number of late jobs. *Annals of Operations Research*, 26(1), 125–133.
- McDiarmid, C. (1989). On the method of bounded differences. *Surveys in Combinatorics*, 141(1), 148–188.
- Pandey, S., Dutta, G., & Joshi, H. (2017). Survey on revenue management in media and broadcasting. *Interfaces*, 47(3), 195–213.
- Phillips, C. A., Uma, R. N., & Wein, J. (2000). Off-line admission control for general scheduling problems. In *Proceedings of the eleventh annual ACM-SIAM symposium on discrete algorithms* (pp. 879–888).

**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.