

1 Nested Active-Time Scheduling

2 Nairen Cao ✉

3 Georgetown University, Washington D.C., USA

4 Jeremy T. Fineman ✉

5 Georgetown University, Washington D.C., USA

6 Shi Li ✉

7 University at Buffalo, Buffalo, New York, USA

8 Julián Mestre ✉

9 The University of Sydney, Sydney, Australia

10 Katina Russell ✉

11 Georgetown University, Washington D.C., USA

12 Seeun William Umboh ✉

13 The University of Sydney, Sydney, Australia

14 — Abstract —

15 The *active-time scheduling problem* considers the problem of scheduling preemptible jobs with
16 windows (release times and deadlines) on a parallel machine that can schedule up to g jobs during
17 each timestep. The goal in the active-time problem is to minimize the number of active steps,
18 i.e., timesteps in which at least one job is scheduled. In this way, the active time models parallel
19 scheduling when there is a fixed cost for turning the machine on at each discrete step.

20 This paper presents a $9/5$ -approximation algorithm for a special case of the active-time scheduling
21 problem in which job windows are laminar (nested). This result improves on the previous best
22 2 -approximation for the general case.

23 **2012 ACM Subject Classification** Theory of computation → Scheduling algorithms

24 **Keywords and phrases** Scheduling algorithms, Active time, Approximation algorithm

25 **Digital Object Identifier** 10.4230/LIPIcs.ISAAC.2022.28

26 **Related Version** *Previous Version*: <https://doi.org/10.1145/3490148.3538554>

27 **1** Introduction

28 The active-time scheduling is the problem of scheduling jobs with windows on a parallel
29 machine so as to minimize the number of timesteps during which machine is on, or active.

30 In the active-time problem [2], we are given as input a set J of n jobs, where each job
31 $j \in J$ has an associated processing time p_j , release time r_j , and deadline $d_j \geq r_j + p_j$,
32 all integers. The jobs are scheduled on a parallel machine that can execute up to g jobs
33 during each step, where g is a positive integer specified as part of the input. The input
34 thus comprises the jobs with their processing times p_j , release times r_j , and deadlines d_j , as
35 well as the machine parameter g , all of which are integers. Time is organized into discrete
36 (integer) steps or slots, and preemption is allowed but only at slot boundaries. We call the
37 time interval $[r_j, d_j)$ the job j 's *window*. Each job j must be fully scheduled within its
38 window, i.e., each job must be assigned to p_j timesteps, where each timestep t to which the
39 job is assigned satisfies $r_j \leq t < d_j$. Moreover, at most g jobs can be scheduled at any step.


40 We say that a timestep t is *active* if the schedule assigns at least one job to step t . The
41 goal in the *active-time scheduling problem* is to find a schedule with minimum number
42 of active steps that schedule all jobs within their windows.



© Nairen Cao, Jeremy T. Fineman, Shi Li, Julián Mestre, Katina Russell and Seeun William Umboh;
licensed under Creative Commons License CC-BY 4.0

33rd International Symposium on Algorithms and Computation (ISAAC 2022).

Editors: Sang Won Bae and Heejin Park; Article No. 28; pp. 28:1–28:17



Leibniz International Proceedings in Informatics

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

43 We assume throughout that the instance is feasible. Testing feasibility (and producing
 44 a schedule) is an easy exercise applying max flow. In fact, this flow-based feasibility test
 45 generalizes to any given subset of active timesteps (see, e.g., Appendix A.1 of [10]). The
 46 active-time scheduling problem thus boils down to figuring out which slots should be activated.

47 Problem History

48 Chang, Gabow, and Khuller [2] introduce the active-time problem and show that the problem
 49 can be solved optimally in polynomial time when the processing times are all one. They also
 50 investigate various generalizations of the problem. For arbitrary integer processing times,
 51 Chang, Khuller, and Mukherjee [3] give two approximation algorithms. First, they give a
 52 rather complex rounding of the natural linear program (LP) that yields a 2-approximation.
 53 They also show that the integrality gap of the natural LP is 2, so the rounding is tight.
 54 Second, they show that, any *minimal feasible solution* yields a 3-approximation. A
 55 minimal feasible solution is a set of active slots such that (i) scheduling the jobs on those slots
 56 is feasible, and (ii) deactivating any slot would render the schedule infeasible. Consequently,
 57 a simple greedy algorithm (choose an arbitrary active slot and deactivate it if the resulting
 58 set of slots is still feasible) is a 3-approximation for the problem. Kumar and Khuller [10]
 59 give a greedy 2-approximation algorithm following the same general strategy of deactivating
 60 slots until reaching a minimal feasible solution, but they choose slots more carefully. They
 61 also exhibit inputs on which their algorithm achieves no better than a $2 - 1/g$ approximation,
 62 so the analysis is effectively tight.

63 A key challenge for improving the approximation ratio for the general active-time problem
 64 is that the natural linear program has an integrality gap of at least $2 - O(1/g)$, which
 65 converges to 2 as $g \rightarrow \infty$. There is also no clear avenue to improve the 2-approximation
 66 obtained by Kumar and Khuller’s [10] greedy approach. Călinescu and Wang [6] suggest
 67 a stronger LP formulation that they conjecture has lower integrality gap, but they only
 68 show that the gap is at least $5/3$ —whether it can lead to better approximations for general
 69 instances remains unknown. Recently, Davies et al.[7] study the active time problem in the
 70 batch scheduling model.

71 Nested active-time scheduling

72 To push past the barriers for the general version of the problem, this paper instead considers
 73 a special case of the active-time problem in which the job windows are laminar (nested).
 74 That is, for each pair of jobs i, j , either the intervals $[r_i, d_i)$ and $[r_j, d_j)$ are disjoint (meaning
 75 either $d_i \leq r_j$ or $d_j \leq r_i$), or one of the intervals is fully contained in the other (i.e., either
 76 $r_i \leq r_j < d_j \leq d_i$ or $r_j \leq r_i < d_i \leq d_j$).

77 Our main result is a $9/5$ -approximation algorithm for the active-time problem with
 78 laminar job windows. Since the simple example exhibiting the integrality gap [3] of 2 for the
 79 natural LP is a nested instance, a different LP formulation is needed. Our algorithm starts
 80 by solving a stronger linear program (LP) for the problem to produce a fractional solution,
 81 then performing a new rounding process over the tree of job windows. The algorithm itself
 82 is not overly complex, but the analysis is not at all straightforward.

83 Restricting our attention to nested windows gives us two advantages. First, assuming
 84 laminar windows allows us to augment the LP to obtain a smaller integrality gap than for the
 85 general case. Notably, our LP includes an additional “ceiling constraint,” which represents a
 86 stronger lower bound on the number of slots required in each window to the volume of jobs
 87 therein. It is not clear how to take advantage of this same constraint in the general version

of the problem. As exhibited by our algorithm the integrality gap of our LP on the nested version of the problem is at most $9/5$, which provides a separation between the nested and general versions of the problem. Secondly, the rounding process itself is inherently tied to the fact that nested windows form a tree. Even ignoring the issue of the larger integrality gap for the general case, it is not clear how to perform a similar rounding for general windows.

In Section 5, we compare our strengthened LP formulation to that proposed by Călinescu and Wang’s [6] and show that both formulations exhibit an integrality gap of at least $3/2$ for nested instances. In Section 6, we show that the nested active time problem is NP complete.

Related work

The objective of the active-time problem is motivated by an application to energy minimization. In this context, the machine can be turned off when no jobs are being executed and it takes the same amount of energy to run regardless of how many jobs are running—but it has a fixed capacity g of how many jobs it can process per active time slot. Energy-aware scheduling is an active area of research [1, 8] that is motivated by the pressing need of modern data centers whose large energy footprint accounts for most of their running costs [13].

There are many variations and generalizations of the basic setup studied in this paper. Below we consider two of its most closely related variants. The reader is referred to the excellent survey by Chau and Li [4] for more related results such as online algorithms for active-time scheduling.

A natural generalization of the basic setup studied in this paper is to have, instead of a single interval, a collection of intervals where each job can be scheduled. Chang *et al.* [2] show that this generalization is NP-hard when $g \geq 3$ even when jobs are unit-length, but that it can be solved in polynomial time when $g = 2$. Furthermore, the problem admits an H_g -approximation for general g via Wolsey’s submodular cover framework [14].

Another related model is the busy-time problem where jobs cannot be preempted and we have parallel machines. This problem is much harder as even testing feasibility for a fixed number of machines is NP-hard. Indeed, the best approximation algorithm for minimizing the number of machines needed when $g = 1$ is the $O\left(\sqrt{\frac{\log n}{\log \log n}}\right)$ -approximation of Chuzhoy *et al.* [5]. Koehler and Khuller [9] show that it is possible minimize the number of machines use and simultaneously achieve a $O(1)$ -approximation on the busy-time objective for instances with uniform processing time; for general instances with arbitrary processing times they can approximate the number of machines by a $O\left(\log \frac{p_{\max}}{p_{\min}}\right)$ factor while keeping the constant approximation bound for the busy-time objective. Liu and Tang [11] studies a generalized busy-time problem on heterogeneous machines and they show an $O(1)$ -approximation algorithm in the offline setting and an $O(\max/\min \text{ job length ratio})$ -competitive algorithm in the online setting.

2 Preliminaries

For an integer p , We use $[p]$ to represent integers from $\{1, 2, \dots, p\}$. Given an instance of the nested active time problem, we define its tree T as follows. Each tree node i is associated with an interval $K(i)$ such that $K(i) = [r_j, d_j)$ for some $j \in J$. If there are several jobs with the same interval, we only create a single tree node. A tree node i' is a child of i if $K(i') \subsetneq K(i)$ and no other node interval is strictly between $K(i)$ and $K(i')$, i.e, there is no node i'' such that $K(i') \subsetneq K(i'') \subsetneq K(i)$. The descendants and ancestors of a node i are denoted $Des(i)$ and $Anc(i)$, respectively. Note that both $Des(i)$ and $Anc(i)$ include i itself.

28:4 Nested Active-Time Scheduling

132 To exclude x , we use $Des^+(i) = Des(i) \setminus \{i\}$ and $Anc^+(i) = Anc(i) \setminus \{i\}$. We define $\text{par}(i)$ to
133 be the parent node of i . W.l.o.g we can assume T is indeed a tree (instead of a forest) since
134 otherwise the instance can be broken into several independent ones.

135 We assume that the tree contains m nodes and each node is associated with an unique id
136 in $[m]$. Now, each job j 's interval is associated with a node in the tree. For a job j , define
137 $k(j)$ to be the tree node i with $K(k(j)) = [r_j, d_j]$; we say j belongs to the node i if $i = k(j)$.
138 For jobs j_1 and j_2 , if $r_{j_1} = r_{j_2}$ and $d_{j_1} = d_{j_2}$, then $k(j_1) = k(j_2)$. Given a node i and a job
139 subset $J' \subseteq J$, $J'(i) = \{j \in J' \mid k(j) = i\}$ is the set of jobs in J' belonging to i . Note that at
140 least one job belongs to each node. Define the length of a node i , which is denoted as $L(i)$,
141 as the $|K(i)| - \sum_{i': \text{par}(i')=i} |K(i')|$, i.e, the number of time slots in the interval $K(i)$, but not
142 in $K(i')$ for any child node i' of i .

143 For simplicity, we use the following shorthand. Given a function or vector f and a set S ,
144 if f outputs reals, then $f(S) = \sum_{e \in S} f(e)$ or $f(S) = \sum_{e \in S} f_e$. If f outputs subsets, then
145 $f(S) = \bigcup_{e \in S} f(e)$ or $f(S) = \bigcup_{e \in S} f_e$.

146 We say that a node i is **rigid** if a feasible solution must open the entire interval $K(i)$.
147 For our rounding algorithm, it will be convenient if the tree is **canonical**.

148 ► **Definition 1** (Canonical trees). *A tree is canonical if it is a binary tree and each leaf node*
149 *is rigid.*

150 First, we transform an arbitrary tree to a binary tree. If a parent node i contains several
151 children nodes i_1, i_2, \dots, i_t , we will create several virtual nodes so that each node contains at
152 most 2 children. Each virtual node's interval is the union of its children's intervals. There
153 are no jobs associated with the virtual nodes and the length of a virtual node i' satisfying
154 $L(i) = 0$. Notice that this transformation adds at most t virtual nodes for a node with t
155 children. In total, this transformation only adds m virtual nodes to a tree that had m nodes
156 originally. In the resulting tree, only internal nodes can be virtual so each leaf node must
157 have at least one job associated with it.

158 We perform one final transformation to make each leaf node rigid. For a leaf node i , let
159 $j \in J(i)$ be a job in i with the longest processing time. If $p_j = L(i)$, then we leave i and the
160 jobs therein unchanged. Otherwise, we can assume that j is scheduled in the first p_j steps of
161 i because j is the longest job in the leaf node and all jobs in the leaf node could choose the
162 leaf's interval to fit in. We transform the instance by creating a virtual child node i' of the
163 leaf i with interval corresponding to the first p_j steps of $K(i)$, and we reduce j 's window to
164 match i' 's. Notice this transformation does not change our solution for the original tree.

165 **3** Algorithm

166 **3.1** Linear Program

167 The linear program is LP (1), which is given in Figure 1(a). In the LP, $x(i)$ denotes the
168 number of time slots opened in node i , excluding times slots in i 's children. $y(i, j)$ denotes
169 the amount of job j that is scheduled in node i . In the LP, OPT_i denotes the smallest
170 number of slots to schedule the jobs in $J(Des(i))$.

171 The objective is to minimize $\sum_{i \in [m]} x(i)$. (2) ensures that every job j is scheduled in at
172 least p_j time slots. (3) ensures that the total number of jobs scheduled in $x(i)$ is at most
173 $g \cdot x(i)$, for each node $i \in [m]$. (4) requires that the number of open time slots in a node
174 $x(i)$ is at most the interval length $L(i)$ of node i . (5) says that we could give at most $x(i)$
175 time slots for a job j . (6) restricts that for each job $j \in J$, j can only be put into nodes in

$$\min \sum_{i \in [m]} x(i) \quad \text{s.t.} \quad (1)$$

$$\sum_{i \in Des(K(j))} y(i, j) \geq p_j, \quad \forall j \quad (2)$$

$$\sum_{j \in J(Anc(i))} y(i, j) \leq g \cdot x(i), \quad \forall i \quad (3)$$


$$x(i) \leq L(i), \quad \forall i \quad (4)$$

$$y(i, j) \leq x(i), \quad \forall i, j \quad (5)$$


$$y(i, j) = 0, \quad \forall i, j \notin J(Anc(i)) \quad (6)$$

$$\sum_{i' \in Des(i)} x(i') \geq 2, \quad \forall i, OPT_i \geq 2 \quad (7)$$

$$\sum_{i' \in Des(i)} x(i') \geq 3, \quad \forall i, OPT_i \geq 3 \quad (8)$$



(b) The open slots from an LP solution



(c) The open slots after performing the LP transformation

(a) Linear program for active time scheduling. By default we restrict $i \in [m]$ and $j \in J$.

Figure 1 (a) is the linear program for active time scheduling. (b) and (c) are an example of a tree before and after running the LP transformation in Lemma 2. The dark slots represent slots have jobs scheduled in them, and the white slots are closed.

176 $Des(K(j))$. (7) and (8) are the key constraints that makes the LP stronger. They are clearly
177 valid; moreover, checking if $OPT_i \geq 2$ ($OPT_i \geq 3$) can be done easily.

178 For simplicity, given a node set $V \subset [m]$ and $J' \subset J$, $y(V, J') = \sum_{i \in V, j \in J'} y(i, j)$; if either
179 V or J' is a singleton, we can replace it by the unique element in it. Let $x(S) = \sum_{i \in S} x(i)$
180 for every $S \subseteq [m]$.

181 After running the LP and getting a solution (x, y) , we will perform a transformation on
182 the solution.

183 3.2 Transformation of LP Solution

184 ► **Lemma 2.** Given a feasible LP solution, we can efficiently output another feasible LP
185 solution such that for any pair of nodes i_1, i_2 such that $i_2 \in Des^+(i_1)$, if $x(i_2) < L(i_2)$, then
186 $x(i_1) = 0$.

187 **Proof.** Suppose there are nodes i_1, i_2 with $i_2 \in Des^+(i_1)$ and $x(i_2) < L(i_2)$ and $x(i_1) > 0$.
188 Then let $\theta = \min\{L(i_2) - x(i_2), x(i_1)\} > 0$. We can move θ fractional open slots from i_1 to i_2 .
189 For every job j , we move $\frac{\theta}{x(i_1)} y(i_1, j)$ fractional assignment of j from i_1 to i_2 . More specifically,
190 let $(x', y') = (x, y)$ initially. We decrease $x'(i_1)$ by θ and increase $x'(i_2)$ by θ . For every
191 $j \in J$, we decrease $y'(i_1, j)$ to $\frac{x'(i_1)}{x(i_1)} y(i_1, j)$ and increase $y'(i_2, j)$ to $y'(i_2, j) + \frac{\theta}{x(i_1)} y(i_1, j)$.
192 Notice that every job j that can be assigned to i_1 can also be assigned to i_2 . It is not hard
193 to show that all constraints remain satisfied by the new solution (x', y') . Finally we update
194 $(x, y) \leftarrow (x', y')$.

195 Notice that after the operation, we have either $x(i_1) = 0$ or $x(i_2) = L(i_2)$. By repeating
196 the procedure a polynomial number of times, we can find a solution (x, y) satisfying the

197 property of the lemma. ◀

198 An example of the LP transformation is shown in Figure 1(b) and 1(c). Lemma 2
 199 implies that for any i with $x(i) > 0$, all of its strict descendants are fully open. We let I be
 200 the set of topmost nodes i with $x(i) > 0$; those are the nodes i with $x(i) > 0$ but all its strict
 201 ancestors i' have $x(i') = 0$.

202 \triangleright **Claim 3.** The following properties hold for I :

203 **(3a)** No node in I is a strict ancestor of another node in I .

204 **(3b)** $Des(I)$ contains all leaves.

205 **(3c)** Every $i \in I$ has $x(i) > 0$.

206 **(3d)** For any $i \in I$ and $i' \in Des^+(i)$, we have $x(i') = L(i')$.

207 **(3e)** For any $i \in I$ and $i' \in Anc^+(i)$, we have $x(i') = 0$.

208 We make one modification to the tree that does not change the instance. For every
 209 $i \in Anc^+(I)$, we can assume that i has exactly two children: if i has only one child, we
 210 remove it from the tree and connect its parent directly to its children. This does not change
 211 the instance since $x(i) = 0$.

212 We also want to mention for any node i such that $x(Des(i)) \in (1, 2)$, i has one child i' ,
 213 which is a leaf with $x(i') = L(i') = 1$ because of the rigidity of the leaf node.

214 3.3 Rounding Algorithm to Obtain an Integral Vector $\tilde{x} \in \{0, 1\}^{[m]}$

215 The rounding algorithm that constructs our integral \tilde{x} is given in Algorithm 1.

■ Algorithm 1 Rounding Algorithm

```

1: let  $\tilde{x}(i) \leftarrow \lfloor x(i) \rfloor, \forall i \in I$  and  $\tilde{x}(i) \leftarrow x(i), \forall i \in [m] \setminus I$ .
2: for every node  $i \in Anc(I)$  from bottom to top
3:   while  $\frac{9x(Des(i))}{5} \geq \tilde{x}(Des(i)) + 1$ 
4:     if  $\exists i' \in Des(i)$  with  $\tilde{x}(i') < x(i')$  then
5:       choose such an  $i'$  arbitrarily
6:       let  $\tilde{x}(i') \leftarrow \lceil x(i') \rceil$ 
7:     else
8:       break

```

216 Clearly, the number of open slots is at most $\frac{9x([m])}{5}$.

217 \blacktriangleright **Lemma 4.** After running the Algorithm 1, $\tilde{x}([m]) \leq \frac{9x([m])}{5}$.

218 4 Feasibility of \tilde{x}

219 In this section, we show that the rounded time slot \tilde{x} is a feasible solution.

220 4.1 A Necessary and Sufficient Condition

221 First we will give an if-and-only-if condition. From now on, for any set $J' \subseteq J$, we define
 222 $J'(i) = J(i) \cap J'$ for every $i \in [m]$ and $J'(S) = J(S) \cap J'$ for every $S \subseteq [m]$.

223 \blacktriangleright **Lemma 5.** Given an integer solution \tilde{x} for the LP, \tilde{x} is feasible if and only if for every
 224 subset $J' \subseteq J$ of jobs, we have

$$225 \sum_{i \in [m]} \min\{|J'(Anc(i))|, g\} \cdot \tilde{x}(i) \geq p(J'). \quad (9)$$

226

227 **Proof.** The only if part is easy to see. Any node i can hold at most $\min\{|J'(Anc(i))|, g\} \cdot$
 228 $\tilde{x}(i)$ volume of jobs in J' . All the jobs in J' should be assigned. If \tilde{x} is feasible, then
 229 $\sum_{i \in [m]} \min\{|J'(Anc(i))|, g\} \cdot \tilde{x}(i) \geq p(J')$.

230 Now we prove the if part, by considering the contra-positive of the statement and applying
 231 the maximum-flow-minimum cut theorem. Assume \tilde{x} is not feasible. We construct a 4-layer
 232 directed network $H = (V_H, E_H)$, where the nodes from left to right are $\{s\}, J, [m]$ and $\{t\}$.
 233 There is an edge from s to every $j \in J$ with capacity p_j , an edge from every $j \in J$ to every
 234 $i \in Des(k(j))$ with capacity $\tilde{x}(i)$, and an edge from every $i \in [m]$ to t with capacity $g \cdot \tilde{x}(i)$.
 235 For a subsets $V' \subseteq V_H$ and a node $v \in V_H \setminus V'$, we use $E_H(V', v)$ to denote the set of edges
 236 from V' to v .

237 As \tilde{x} is not feasible, there is a s - t cut in the network with capacity less than $p(J)$.
 238 Let (A, B) be the cut: $s \in A, t \in B$ and $A \uplus B = V_H$. Its cut value, which is $p(B \cap$
 239 $J) + \sum_{i \in B \cap [m]} |E_H(A \cap J, i)| \cdot \tilde{x}(i) + g \cdot \tilde{x}(A \cap [m])$, is less than $p(J)$. This is equivalent to
 240 $\sum_{i \in B \cap [m]} |E_H(A \cap J, i)| \cdot \tilde{x}(i) + g \cdot \tilde{x}(A \cap [m]) < p(A \cap J)$. Let $J' = A \cap J$. Then the contribution
 241 of a node $i \in [m]$ to the left-side is either $|E_H(J', i)| \cdot \tilde{x}(i)$ (if $i \in B$), or $g \cdot \tilde{x}(i)$ (if $i \in A$),
 242 which is lower bounded by $\min\{|E_H(J', i)|, g\} \cdot \tilde{x}(i)$. Noticing $|E_H(J', i)| = |J'(Anc(i))|$, we
 243 have $\sum_{i \in [m]} \min\{|J'(Anc(i))|, g\} \cdot \tilde{x}(i) < p(J')$. This finishes the proof of the if part. ◀

244 ▶ **Lemma 6.** *In Lemma 5, it is sufficient to consider the sets $J' \subseteq J$ satisfying the following*
 245 *property:*

246 **(6a)** $p_j > \tilde{x}\left(\{i \in Des(K(j)) : |J'(Anc(i))| \leq g\}\right), \forall j \in J'$.

247 **Proof.** Suppose J' does not satisfy the property. Then for some $j \in J'$ we have $\tilde{x}\left(\{i \in$
 248 $Des(K(j)) : |J'(Anc(i))| \leq g\}\right) \geq p_j$. Then removing j from J' will decrease the left side of
 249 (9) by $\tilde{x}\left(\{i \in Des(K(j)) : |J'(Anc(i))| \leq g\}\right)$, and the right side by p_j . Thus, the inequality
 250 (9) for J' will be implied by the inequality for $J' \setminus \{j\}$. ◀

251 Once we have the if-and-only-if condition for the feasibility, the main lemma we need to
 252 prove is the following:

253 ▶ **Theorem 7.** *For every subset $J' \subseteq J$ of jobs satisfying the property in Lemma 6, we have*
 254 *(9).*

255 The rest of the section is dedicated to the proof of Theorem 7. From now on we fix a
 256 subset $J' \subseteq J$ satisfying the property of Lemma 6. Our goal is to prove (9).

257 For notational convenience, let $u_i = \min\{|J'(Anc(i))|, g\}$ and $w_i = u_i x(i)$ and $\tilde{w}_i = u_i \tilde{x}(i)$
 258 for every $i \in [m]$. Thus, (9) is simply written as $p(J') \leq \tilde{w}([m])$. Recall that $y(V, J') =$
 259 $\sum_{i \in V, j \in J'} y(i, j)$ for a given $V \subseteq [m], J' \subseteq J$. We have $p(J') = y([m], J') = y(Des(I), J')$
 260 and $\tilde{w}([m]) = \tilde{w}(Des(I))$. Thus, we need to prove

261 $y(Des(I), J') \leq \tilde{w}(Des(I)),$ (10)
 262

263 for every $J' \subseteq J$ satisfying the property in Lemma 6.

264 The following simple claim will be used multiple times:

265 ▷ **Claim 8.** For every $i \in [m]$, we have $y(i, J') \leq w_i = u_i x(i)$.

266 **Proof.** If $u_i = g$, we use (3) in the LP. If $u_i < g$, then we use (5) and (6). ◀

267 **4.2 Construction of Triples**

268 For nodes $i \in [m] \setminus I$, we have $\tilde{x}(i) = x(i)$. For nodes $i \in I$ with $x(Des(i)) \notin (1, \frac{10}{9})$, we
 269 have $\tilde{x}(i) = \lceil x(i) \rceil$ since $\frac{9x(Des(i))}{5} \geq \lceil x(Des(i)) \rceil$. Thus, for these nodes i , Claim 8 implies
 270 $y(i, J') \leq \tilde{w}_i$. The critical nodes are those $i \in I$ with $x(Des(i)) \in (1, \frac{10}{9})$.

271 With this in mind, we classify nodes in I into two types: a node $i \in I$ is of

272 ■ type-B if $x(Des(i)) \in \{1\} \cup [\frac{4}{3}, \infty)$, and

273 ■ type-C if $x(Des(i)) \in (1, \frac{4}{3})$.

274 In the definition, we use $\frac{4}{3}$ instead of $\frac{10}{9}$ to create some buffers. Furthermore, a type-C node
 275 $i \in I$ is of

276 ■ type-C₁ if $\tilde{x}(Des(i)) = 1$, and

277 ■ type-C₂ if $\tilde{x}(Des(i)) = 2$.

278 **At most 2 type-C nodes** Before going to the triples, we first solve the case at most 2
 279 type-C nodes are in I . At the same time, if 1 type-B node exists, then all type-C nodes are
 280 type-C₂.

281 ► **Lemma 9.** *If there are at most 2 type-C nodes and at least 1 type-B node in I , then all*
 282 *type-C are type-C₂.*

283 **Proof.** Let i_1 and i_2 (if exists) be the type-C node and i_3 be the type-B node. Notice that
 284 $\frac{9}{5}x(i_3) - \lceil x(i_3) \rceil \geq 0.4$. We have $\frac{9}{5}(x(i_1) + x(i_3)) \geq x(i_1) + 0.8 + \lceil x(i_3) \rceil + 0.4 \geq \lceil x(i_1) \rceil + \lceil x(i_3) \rceil$
 285 and $\frac{9}{5}(x(i_1) + x(i_2) + x(i_3)) \geq x(i_1) + 0.8 + x(i_2) + 0.8 + \lceil x(i_3) \rceil + 0.4 \geq \lceil x(i_1) \rceil + \lceil x(i_2) \rceil + \lceil x(i_3) \rceil$.
 286 In either case, algorithm 1 can afford to round up all type-C nodes. ◀

287 Based on Lemma 9, if at least one type-B node is in I , then we already rounded up all
 288 type-C nodes. Assume that there is no type-B node in I , then we have $x([m]) \leq 2 \times 4/3 = 8/3$.
 289 Due to the LP constraint (7) and (8), $x(i)$ are integer for all $i \in [m]$ and this contradicts to
 290 the assumption that no type-B node is in I .

291 **More than 2 type-C nodes** When we have at least 3 type-C nodes in I , we want to
 292 create disjoint triples of type-C nodes. Each triple contains 1 type-C₁ node, and 2 type-C₂
 293 nodes. Moreover, all the type-C₁ nodes are contained in these triples. Later for each triple
 294 (i_1, i_2, i_3) we constructed, we shall prove $y(Des(\{i_1, i_2, i_3\}), J') \leq \tilde{w}(Des(\{i_1, i_2, i_3\}))$. This
 295 will prove (10).

296 The construction of triples are given in Algorithm 2. Notice that this is not a part of
 297 our algorithm for solving the active time scheduling problem; it is only used in the analysis.
 298 If we have a type-C₁ node i_1 and a type-C₂ node i_2 as brothers, then we say (i_1, i_2) is a
 299 C₁C₂-brother-pair. In our triples, we make sure that we do not break C₁C₂-brother-pairs:
 300 for such a pair (i_1, i_2) , there must be some C₂-node i_3 such that (i_1, i_2, i_3) is a triple we
 301 constructed.

■ **Algorithm 2** Construction of Triples

-
- 1: $\text{triples} \leftarrow \emptyset$, set all type-C₁ nodes as uncovered, and all type-C₂ nodes as unused.
 - 2: **for** every node $i \in Anc(I)$ with $|Des(i) \cap I| \geq 3$ from bottom to top
 - 3: **while** \exists an uncovered type-C₁ node $i_1 \in Des(i)$
 - 4: choose two unused type-C₂ nodes $i_2, i_3 \in Des(i)$, without breaking C₁C₂-brother-
 pairs ▷ See Lemma 10 and
 - Lemma 11
 - 5: add (i_1, i_2, i_3) to triples , claim i_1 is covered, and i_2 and i_3 are used.
-

302 ► **Lemma 10.** *In Step 4 of Algorithm 2, there are at least two unused type-C₂ nodes in*
 303 *Des(i).*

304 **Proof.** Let n_1 and n_2 be the number of type-C₁ and type-C₂ nodes in $Des(i) \cap I$ respectively.
 305 Let $n' = n_1 + n_2$. We shall prove $n_2 \geq 2n_1$, which implies that we will not run out of type-C₂
 306 nodes and proves the lemma.

307 First, we consider the case where there is no type-B node in $Des(i) \cap I$. Then $n' =$
 308 $|Des(i) \cap I| \geq 3$. Moreover, $n_1 + 2n_2 \geq \lfloor \frac{9n'}{5} \rfloor$ or $n_1 = 0$ by our rounding algorithm in
 309 Algorithm 1. In the latter case we clearly have $n_2 \geq 2n_1$. So, we assume the former. Then
 310 $n_2 \geq \lfloor \frac{4n'}{5} \rfloor$. Notice that $\lfloor \frac{4n'}{5} \rfloor \geq \frac{4n'}{5} - \frac{4}{5}$ and thus $\lfloor \frac{4n'}{5} \rfloor \geq \frac{2n'}{3}$ whenever $n' \geq 6$. One can
 311 check that when $n' \in \{3, 4, 5\}$ we have $\lfloor \frac{4n'}{5} \rfloor \geq \frac{2n'}{3}$. Therefore if $n' \geq 3$, we have $\lfloor \frac{4n'}{5} \rfloor \geq \frac{2n'}{3}$.
 312 So $n_2 \geq 2n_1$.

313 Now we consider the other case where there is at least one type-B node in $Des(i) \cap I$. If
 314 $n_1 = 0$, then $n_2 \geq 2n_1$ and thus we assume $n_1 > 0$. Then we have $n_1 + 2n_2 \geq \lfloor \frac{9n'}{5} + \frac{2}{5} \rfloor$, as
 315 the type B node i^* have $\frac{9x(Des(i^*))}{5} \geq \lceil x(Des(i^*)) \rceil + \frac{2}{5}$. This implies $n_2 \geq \lfloor \frac{4n'}{5} + \frac{2}{5} \rfloor$. Then
 316 $n_2 \geq \frac{2n'}{3}$ as $\lfloor \frac{4n'}{5} + \frac{2}{5} \rfloor \geq \frac{2n'}{3}$ for every integer $n' \geq 0$. Thus, $n_2 \geq 2n_1$. ◀

317 4.3 Proof of (10) Using the Triples

318 In this section, we prove (10) by using the constructed triples. First, we show that they
 319 satisfy some good properties:

320 ► **Lemma 11.** *For every $(i_1, i_2, i_3) \in \text{triples}$, one of the following two conditions hold:*

321 **(11a)** $i_2, i_3 \in Des^+(\text{par}(i_1))$.

322 **(11b)** i_1 and i_2 are brothers, and $i_3 \in Des^+(\text{par}(\text{par}(i_1)))$.

323 **Proof.** Consider any type-C₁ node i_1 . Let i^* and i' be the parent and brother of i_1 respect-
 324 ively.

325 First consider the case that $i' \notin I$. By (3a) and (3b), we have $|Des(i^*) \cap I| \geq 3$. By
 326 Algorithm 2 and Lemma 10, i' will be covered when we are at the iteration $i = i^*$ in
 327 Algorithm 2. So, the first property holds.

328 Then consider the other case that $i' \in I$. Then i' can not be of type-B since otherwise
 329 i_1 should be of type-C₂. i' can not be of type-C₁ since we could open 3 slots in $Des(i^*)$.
 330 Therefore i' must be of type-C₂ and (i_1, i') is a C₁C₂-brother-pair. In this case the second
 331 property holds. ◀

332 See Figure 2 for the two cases obtained in Lemma 11, which will be used again in the
 333 proof of the following Lemma:


334 ► **Lemma 12.** *For every $(i_1, i_2, i_3) \in \text{triples}$, we have*

$$335 \quad y(Des(\{i_1, i_2, i_3\}), J') \leq \tilde{w}(Des(\{i_1, i_2, i_3\})).$$

337 **Proof.** Recall that i_1 is of type C₁ and i_2 and i_3 are of type-C₂. Let i'_1, i'_2 and i'_3 be the
 338 children of i_1, i_2 and i_3 respectively. By Lemma 11, either a or b holds.

339 We first assume a. Let $i = \text{par}(i_1)$. See Figure 2(a) for all illustration of nodes used in
 340 this case.

28:10 Nested Active-Time Scheduling



■ **Figure 2** The two cases in Lemma 11 and 12.

341 By (7) in the LP, $J(Des(i_1))$ can be scheduled in the one slot in i'_1 since $x(Des(i_1)) < 2$.
 342 That is, all the jobs in the set has size 1 and there are at most g of them. So, we have

$$343 \quad y(Des(i_1), J'(Des(i_1))) \leq |J'(Des(i_1))| \leq u_{i'_1} x(i'_1). \quad (11)$$

$$344 \quad y(Des(i_1), J'(Anc^+(i_1))) \leq \frac{10}{9} \min\{|J'(Anc^+(i_1))|, g\}$$

$$345 \quad \leq (\tilde{x}(i_2) - x(i_2) + \tilde{x}(i_3) - x(i_3)) \min\{|J'(Anc^+(i_1))|, g\}$$

$$346 \quad \leq (\tilde{x}(i_2) - x(i_2))u_{i_2} + (\tilde{x}(i_3) - x(i_3))u_{i_3} \quad (12)$$

$$347 \quad y(Des(\{i_2, i_3\}), J') \leq u_{i_2} x(i_2) + u_{i'_2} x(i'_2) + u_{i_3} x(i_3) + u_{i'_3} x(i'_3) \quad (13)$$

349 The first inequality of (11) is by that all jobs in $J'(Des(i_1))$ have size 1, and the second
 350 one is by $u_{i'_1} = \min\{|J'(Anc^+(i_1))|, g\} \geq |J'(Des(i_1))|$ and $x(i'_1) = 1$. The first inequality
 351 of (12) is by that $x(Des(i_1)) < \frac{10}{9}$. The second one follows from $\tilde{x}(i_2) - x(i_2) \geq \frac{2}{3}$ and
 352 $\tilde{x}(i_3) - x(i_3) \geq \frac{2}{3}$. The last one used that every job in $J'(Anc^+(i_1))$ can be assigned to i_2
 353 and i_3 . (13) is by Claim 8.

354 Adding (11), (12) and (13), we have

$$355 \quad y(Des(\{i_1, i_2, i_3\}), J') \leq u_{i'_1} x(i'_1) + u_{i_2} \tilde{x}(i_2) + u_{i'_2} x(i'_2) + u_{i_3} \tilde{x}(i_3) + u_{i'_3} x(i'_3)$$

$$356 \quad = \tilde{w}(Des(\{i_1, i_2, i_3\}))$$

358 Now we consider the case that b holds. Let $i = \text{par}(i_1) = \text{par}(i_2)$. See Figure 2(b) for
 359 illustration of the nodes.

360 First, if $u_{i_2} = g$, then we have $(\tilde{x}(i_2) - x(i_2))u_{i_2} \geq (x(i_1) - \tilde{x}(i_1))u_{i_1}$ as $\tilde{x}(i_2) - x(i_2) \geq \frac{2}{3}$
 361 and $x(i_1) - \tilde{x}(i_1) < \frac{1}{9}$. This is $\tilde{w}(i_1) + \tilde{w}(i_2) \geq w(i_1) + w(i_2)$. Thus, $y(Des(\{i_1, i_2, i_3\}), J') \leq$
 362 $w(Des(\{i_1, i_2, i_3\})) \leq \tilde{w}(Des(\{i_1, i_2, i_3\}))$ as $\forall i' \in \{i_3, i'_1, i'_2, i'_3\}$ we have $w(i') \leq \tilde{w}(i')$.

363 So assume $u_{i_2} < g$. As we assumed J' satisfies (6a), every job $j \in J'(Anc^+(i_2))$ has $p_j > 1$.
 364 All jobs in $J(i)$ have $p_j \leq 2$ since otherwise we would have $x(Des(i)) \geq 3$ by LP constraint
 365 (2) and (5). Also all jobs in $J(i_2)$ have size 1. So all jobs in $J'(i)$ have size 2, and $J'(i_2) = \emptyset$.

366 Then we have

$$367 \quad y(\text{Des}(i_1), J'(\text{Des}(i_1))) + |J'(i)| \leq |J'(\text{Des}(i_1))| + |J'(i)| \leq u_{i'_1} x(i'_1). \quad (14)$$

$$368 \quad y(\text{Des}(i_1), J'(\text{Anc}^+(i))) \leq \frac{10}{9} |J'(\text{Anc}^+(i))| \leq (\tilde{x}(i_2) - x(i_2) + \tilde{x}(i_3) - x(i_3)) |J'(\text{Anc}^+(i))| \\ 369 \quad \leq (\tilde{x}(i_2) - x(i_2)) |J'(\text{Anc}^+(i))| + (\tilde{x}(i_3) - x(i_3)) u_{i_3} \quad (15)$$

$$370 \quad y(\text{Des}(\{i_2, i_3\}), J' \setminus J'(i)) + |J'(i)| \leq |J'(\text{Anc}^+(i))| \cdot x(i_2) + u_{i'_2} x(i'_2) + u_{i_3} x(i_3) \\ 371 \quad + u_{i'_3} x(i'_3) + |J'(i)| \cdot \tilde{x}(i_2) \quad (16)$$

373 The first inequality of (14) all jobs in $J'(\text{Des}(i_1))$ have size 1, and the second inequality
374 is by that $|J'(\text{Des}(i_1))| + |J'(i)| \leq g$ since otherwise $\text{OPT}_i \geq 3$. The proof of (15) is similar
375 to that of (12). Notice that every job in $J'(\text{Anc}^+(i))$ can be assigned to i_2 and i_3 . The
376 inequality in (16) used Claim 8 for i'_2, i_3 and i'_3 .

377 Adding (14), (15) and (16), we get

$$378 \quad y(\text{Des}(i_1, i_2, i_3), J' \setminus J'(i)) + 2|J'(i)| \\ 379 \quad \leq u_{i'_1} x(i'_1) + (|J'(\text{Anc}^+(i))| + |J'(i)|) \tilde{x}(i_2) + u_{i'_2} x(i'_2) + u_{i_3} \tilde{x}(i_3) + u_{i'_3} x(i'_3) \\ 380 \quad = u_{i'_1} x(i'_1) + u_{i_2} \tilde{x}(i_2) + u_{i'_2} x(i'_2) + u_{i_3} \tilde{x}(i_3) + u_{i'_3} x(i'_3) = \tilde{w}(\text{Des}(i_1, i_2, i_3)).$$

382 Notice that $y(\text{Des}(i_1, i_2, i_3), J'(i)) = p(J'(i)) = 2|J'(i)|$, we have $y(\text{Des}(i_1, i_2, i_3), J') \leq$
383 $\tilde{w}(\text{Des}(i_1, i_2, i_3))$. \blacktriangleleft

384 With Lemma 12, we can prove (10). First $\sum_{i \in *} y(\text{Des}(i), J') \leq \sum_{i \in *} \tilde{w}(\text{Des}(i))$, where
385 $i \in *$ is over all nodes in the triples we constructed. For all the other nodes $i \in I$, we have
386 $y(\text{Des}(i), J') \leq w(\text{Des}(i)) \leq \tilde{w}(\text{Des}(i))$. Therefore we have $y(\text{Des}(I), J') \leq \tilde{w}(\text{Des}(I))$, which
387 is exactly (10). Combining with Lemma 4, we have following theorem,

388 **► Theorem 13.** *There exists a 1.8-approximation polynomial-time algorithm for the nested*
389 *active-time problem.*

390 5 Integrality gap

391 For the general (non-nested) version of the active scheduling problem, Călinescu and Wang [6]
392 proposed a slightly stronger than our LP from Figure 1a and showed a non-nested instance
393 where the integrality gap approaches $5/3$ as $g \rightarrow \infty$. In this section, we show that their LP
394 and our LP have an integrality gap of at least $3/2$ on nested instances.

395 To define their LP, we need some notation. Let $\mathcal{T} = [\min_{j \in J} r_j, \max_{j \in J} d_j]$ denote the
396 set of time steps between the earliest release time and the latest deadline. For an interval of
397 time $I = [t_1, t_2)$ for some $t_1, t_2 \in \mathcal{T}$ and a job j , let $q_j(I)$ denote the minimum number of
398 slots within I that job j needs to occupy in a feasible solution even if all slots outside of I
399 were active and available to j . The variable $x(t)$ denotes the extent to which the slot t is
400 active and $y(t, j)$ denotes the extent to which job j is assigned to slot t . See Figure 3.

401 **► Lemma 14.** *The linear program in Figure 3 has an integrality gap of at least $3/2$ on nested*
402 *instances.*

403 **Proof.** The integrality gap instance consists of one long job j_0 with processing time g and
404 window $[0, 2g)$, and g groups of g jobs. For $0 \leq i < g$, the i -th group consists of g jobs with
405 unit processing time and window $[2i, 2i + 2)$.

406 Consider the following LP solution (x, y) : open each slot $t \in \mathcal{T} = [0, 2g)$ to an extent
407 of $x(t) = (g + 2)/2g$ for a total of $g + 2$. For each $0 \leq i < g$, the LP schedules $1/2$

$$\begin{array}{l}
 \min \sum_{t \in \mathcal{T}} x(t) \\
 \text{s.t.} \quad \sum_{t \in [r_j, d_j)} y(t, j) \geq p_j \quad \forall j \in J \\
 \sum_{j \in J} y(t, j) \leq g \cdot x(t) \quad \forall t \in \mathcal{T} \\
 y(t, j) \leq x(t) \quad \forall t \in \mathcal{T}, \forall j \in J \\
 x(t) \leq 1 \quad \forall t \in \mathcal{T} \\
 \sum_{t \in [t_1, t_2)} x(t) \geq \left\lceil \frac{\sum_{j \in J} q_j(I)}{g} \right\rceil \quad I = [t_1, t_2), \forall t_1 \in \mathcal{T}, \forall t_2 \in \mathcal{T}
 \end{array}$$

■ **Figure 3** Călinescu and Wang’s linear program for active time scheduling [6]

408 unit of j_0 and $1/2$ unit of each job j in the i -th group in slots $2i$ and $2i + 1$; that is,
 409 $y(2i, j_0) = y(2i + 1, j_0) = 1/2$ and $y(2i, j) = y(2i + 1, j) = 1/2$ for each job j in group i .

410 We now argue that (x, y) satisfies the ceiling constraints of [6]’s LP; it is easy to check
 411 that the other constraints are satisfied. Consider an interval I . Since the long job j_0 ’s window
 412 is $[0, 2g)$ and j_0 has length g , we have that $q_{j_0}(I) = 0$ if $|I| \leq g$ and $q_{j_0}(I) = |I| - g$ if $|I| > g$.
 413 This is because there are $2g - |I|$ slots outside of I . For a job j in group i , since it has unit
 414 length, we have $q_j(I) = 1$ if I contains its window $[2i, 2i + 2)$ and $q_j(I) = 0$ otherwise.

415 Combining the above, the LP constraint on interval I is

$$416 \quad \sum_{t \in I} x(t) \geq \left\lceil \frac{\max\{0, |I| - g\} + g|\{i \mid I \supseteq [2i, 2i + 2)\}|}{g} \right\rceil.$$

417 The tightest constraints are when I is the union of windows of consecutive groups. Thus, it
 418 suffices to argue that these constraints are satisfied. Suppose $I = [2i', 2(i' + k - 1) + 2)$, i.e.,
 419 it is the union of windows of k consecutive groups. Note that $|I| = 2k$

420 If $k \leq g/2$, then the LP constraint on I is $\sum_{t \in I} x(t) \geq k$. This is satisfied since
 421 $x_{2i} + x_{2i+1} = (g + 2)/g$ for each group i . If $k > g/2$, then the LP constraint on I is
 422 $\sum_{t \in I} x(t) \geq 1 + k$. This is also satisfied since $\sum_{t \in I} x(t) = k(g + 2)/g = k + 2k/g > k + 1$.
 423 Thus, (x, y) is a feasible solution to the LP that opens $g + 2$ slots fractionally.

424 We claim that any integral solution (x', y') needs to open at least $3g/2$ slots. Let k be
 425 the number of groups i such that y' schedules at least one unit of the long job j_0 in the
 426 window $[2i, 2i + 2)$. Consider the window $[2i, 2i + 2)$. Since there are g unit jobs that need
 427 to be scheduled in $[2i, 2i + 2)$, x' opens two slots in the window if y' schedules j_0 in the
 428 window and only one slot otherwise. Thus, y' opens exactly $g + k$ slots. Since each window
 429 $[2i, 2i + 2)$ has only two slots, and j_0 has length g , we have that $k \geq g/2$. Therefore, any
 430 integral solution needs to open at least $3g/2$ slots. Thus, the integrality gap of the LP is at
 431 least $\frac{3g}{2(g+2)}$ which converges to $3/2$ for large g . ◀

432 6 NP-COMPLETENESS

433 In this section, we show that the decision version of the nested active time problem is NP
 434 complete. Very recently, Sagnik and Manish [12] showed the general case is NP complete.

435 Unfortunately, their proof uses crossing intervals (i.e., intervals that overlap but neither is
 436 included in the other). Our proof reduces the nested active time problem to a new problem
 437 that we call **prefix sum cover**, which is related to the classic set cover problem.

438 **Prefix sum cover problem.**

439 For any pair of d -dimensional vectors $v = (v_1, v_2, \dots, v_d), w = (w_1, w_2, \dots, w_d) \in R^d$, we say
 440 $v \prec w$ if and only if for all $j \in [1, d]$, $\sum_{i \leq j} v_i \geq \sum_{i \leq j} w_i$. In the prefix sum cover problem,
 441 we are given n vectors $u_1, u_2, \dots, u_n \in N_+^d$, a target vector $v \in N^d$ and an integer number k ,
 442 and we want to find k vectors $u_{l_1}, u_{l_2}, \dots, u_{l_k}$ such that $\sum_{i \leq k} u_{l_i} \prec v$.

443 Moreover, for the purposes of our reduction, we consider a restricted version of the
 444 problem. Let W be the maximum scalar that appears in any of the vectors u_1, \dots, u_n and
 445 v . First, we require that both d and W be bounded by some polynomial of n . For a vector
 446 $w \in N^d$, let $[w]_j$ be its j -th dimension value. Second, for each $i \in [1, n]$, we require that
 447 $[u_i]_1 \geq [u_i]_2 \geq \dots \geq [u_i]_d$, and $[v]_1 \geq [v]_2 \geq \dots \geq [v]_d$, i.e., all vectors are non-decreasing.
 448 Lastly, we require that all vectors are non-negative and integral.

449 We show in Appendix 6 that prefix sum cover is NP complete even under these restriction.

450 **Remark:** Notice that the prefix sum cover problem is almost the same as set cover
 451 problem except for the "order" relation. We can think of the set cover problem as requiring
 452 that each dimension of the sum vector is greater than the target vector, while in the prefix
 453 sum cover problem, the requirement is "prefix sum".

454 **Reduction**

455 Now we will reduce the prefix sum cover problem to the active time problem. Let $(\{u_1, u_2,$
 456 $\dots, u_n\}, v, k)$ be the prefix sum cover instance. Our nested active time instance is defined by
 457 a set of jobs J and it uses $p = dW$ machines. Our instance is made up of three kinds of jobs:

- 458 ■ For each vector u_i , and each $w \in [2, W]$, we have $p - |\{j \in [1, d] \mid [u_i]_j \geq w\}|$ rigid unit
 459 length jobs, each with window consisting of a single slot $[(i-1)W + w - 1, (i-1)W + w]$.
- 460 ■ For each vector u_i , we also have $\sum_{j \leq d} [u_i]_j - d$ flexible unit jobs with window $[(i-1)W, iW]$.
- 461 ■ Finally, we have jobs that depend on the target vector. For each $j \in [1, d]$, we have a job
 462 with length $[v]_j$ and window $[0, nW]$.

463 We denote each of these sets of job with S_1 (rigid jobs), S_2 (flexible jobs associate with each
 464 u_i vector), and S_3 (jobs associated with the target vector).

465 Let us try to schedule this instance, starting with S_1 . Since the jobs in S_1 are rigid, we
 466 must open all slots in $[(i-1)W + 1, iW]$. Notice that each of these slots has at least $p - d$
 467 jobs in S_1 , so each of these time slots has at most d unused machines after scheduling S_1 .

468 Next we will try to fit jobs from S_2 into $[(i-1)W, iW]$. Observe that the total capacity
 469 in the window $[(i-1)W + 1, iW]$ is $p(W-1)$ and that the jobs from S_1 take up up
 470 $\sum_{w \in [2, W]} (p - |\{j \in [1, d] \mid [u_i]_j \geq w\}|)$ capacity. Further observe that

$$471 \sum_{w \in [2, W]} (p - |\{j \in [1, d] \mid [u_i]_j \geq w\}|) + \sum_{j \leq d} [u_i]_j - d = p(W-1).$$

472 Therefore, if we do not open the slot $[(i-1)W, (i-1)W + 1]$, then the jobs from S_1 and S_2
 473 will use up all of the available capacity in the time window $[(i-1)W, iW]$. This is important,
 474 as it means that we cannot schedule any job from S_3 in this window.

475 We say that the time slots $[(i-1)W, (i-1)W + 1]$ for $i \in [n]$ are *special*. Since all
 476 non-special slots in $[0, nW]$ must be open, the problem boils down to opening as few special
 477 slots as possible to accommodate the jobs in S_3 .

28:14 Nested Active-Time Scheduling

478 Suppose we open the special time slot $[(i-1)W, (i-1)W+1]$. We claim that all jobs in
 479 S_2 will be assigned to the special time slot. Indeed, even after all S_2 jobs are assigned to
 480 this slot, there are still $p - (\sum_{j \leq d} [u_i]_j - d) \geq d$ unused machines in it, while we can only
 481 have at most d unused machines in each time slots in $[(i-1)W+1, iW]$ after scheduling S_1 .

482 Let *configuration* be a sequence (z_1, z_2, \dots, z_M) , where z_i is the number of machines unused
 483 in time slot $[i-1, i]$. Thus, once we have chosen which special slots to open, we get the
 484 configuration which tells us how many machines are left unused in each time slot. In the
 485 remainder of this section, we give an if-and-only-if condition on whether a configuration can
 486 fit all jobs from S_3 .

487 Assume the machines are numbered from 1 to p . For any given configuration, we can
 488 assume without loss of generality that if we have z_t unused capacity at time slot $[t-1, t]$
 489 then machines 1 through z_t are unused; i.e, we always leaves smaller index machine unused.
 490 Let e_j be the number of empty time slots at machine j . We give an if-and-only-if condition
 491 for the feasibility based on the e_j values.

492 ► **Lemma 15.** *Given a configuration, let e_j be the machine unused slot defined above and J'
 493 be a set of $q \leq p$ jobs with no release time and due time constraint. Let $l_1 \geq l_2 \geq \dots \geq l_q$
 494 be the lengths of the jobs in J' . The configuration can fit all jobs in J' if and only if
 495 $\sum_{i \leq j} e_i \geq \sum_{i \leq j} l_i$ for all $j \in [1, q]$.*

496 **Proof.** To identify each job, when we say the i -th job, we refer the job with length l_i .

497 **If part** Suppose $\sum_{i \leq j} e_i \geq \sum_{i \leq j} l_i$ for all $j \in [1, q]$. Now, we prove the following
 498 statement by induction on k : if $\sum_{i \leq k} e_i \geq \sum_{i \leq k} l_i$ for any k , then we can fit jobs l_1, l_2, \dots, l_k
 499 into the first k machines. The base case is $k = 1$: since $e_1 \geq l_1$, then we can fit l_1 into the
 500 first machine. Now we prove the inductive case. Suppose we can fit the first k jobs into the
 501 first k machines and we want to fit the first $k+1$ jobs into the first $k+1$ machines. Now,
 502 we first try to fit the first job to the first $k+1$ machines and then use our induction. We
 503 fit the first job in following sequence: we first use machine $k+1$ and if we cannot fit all
 504 of the first job in it, we use the machine k and repeat this process, i.e, fit the first job by
 505 using machine with decreasing index. Notice that we use at most e_1 time slots to fit the
 506 first job, thus the above process will finally fit the first job inside but use some time slot of
 507 machines from $k+1$ to 1. Now, let $e'_1, e'_2, e'_3, \dots, e'_{k+1}$ be the machine unused time slot after
 508 we fit the first job inside. We know that $e'_j \geq e_{j+1}$ for $j \in [1, k]$ since a job cannot use the
 509 same time slot twice. Another point is if $e'_j < e_j$, i.e, we use some time slot of machine j ,
 510 then $\sum_{i > j} (e_i - e'_i) = e_{j+1}$. Now, notice that we have jobs l_2, l_3, \dots, l_{k+1} , and the new time
 511 slot is e'_1, e'_2, \dots, e'_k , if we can show $\sum_{i \leq j} e'_i \geq \sum_{i \leq j} l_{i+1}$ for all $j \leq k$, then by induction, we
 512 can fit the second to the $(k+1)$ -th jobs to the first k machines. If $e'_j = e_j$, then we can say
 513 $e'_{j'} = e_{j'}$ for all $j' \leq j$, and thus $\sum_{i \leq j} e'_i = \sum_{i \leq j} e_i \geq \sum_{i \leq j} l_i \geq \sum_{i \leq j} l_{i+1}$. If $e'_j < e_j$, we
 514 have $l_1 = \sum_{i \leq k} (e_i - e'_i) = e_{j+1} + \sum_{i \leq j} (e_i - e'_i)$, thus

$$\begin{aligned} \sum_{i \leq j+1} e_i &\geq \sum_{i \leq j+1} l_i = l_1 + \sum_{i \leq j} l_{i+1} = e_{j+1} + \sum_{i \leq j} (e_i - e'_i) + \sum_{i \leq j} l_{i+1} \\ &\Rightarrow \sum_{i \leq j} e_i - \sum_{i \leq j} (e_i - e'_i) \geq \sum_{i \leq j} l_{i+1} \Rightarrow \sum_{i \leq j} e'_i \geq \sum_{i \leq j} l_{i+1} \end{aligned}$$

516 In either case, we know that the remaining jobs could be fitted into machines from 1 to
 517 k , thus we could fit all jobs if for all $j \in [1, p]$, $\sum_{i \leq j} e_j \geq \sum_{i \leq j} l_j$.

518 **Only if part** If for some $j \in [1, q]$, $\sum_{i \leq j} e_i < \sum_{i \leq j} l_i$, then the configuration is impossible
 519 to fit jobs l_1, l_2, \dots, l_j . Consider a time slot $[t-1, t]$, when we fit a job inside, we always use

520 the machine with smallest index, if we could fit l_1, l_2, \dots, l_j into the time slot, then we will
 521 use machines with index at most j . If at any time slot, we use a machine with index greater
 522 than j , then we know we already use all machines from 1 to j , however, we have only j jobs
 523 now. Thus, if we could fit the first j jobs into the time slot, we can use the first j machines
 524 to fit those jobs. However, for the first j machines, the total capacity $\sum_{i \leq j} e_i$ is less than
 525 the length of all jobs, i.e., $\sum_{i \leq j} l_i$. Therefore, it is not possible to fit the first j jobs into the
 526 configuration. ◀

527 Now we show how to apply it to the active time instance. We will set $J = S_3$ and
 528 $q = d$. For any interval $[(i-1)W, iW]$, let $e_{1,i}, e_{2,i}, \dots, e_{d,i}$ be the unused time slot for
 529 machine 1 to d in this interval. If we close the special time slot $[(i-1)W, (i-1)W+1]$,
 530 then there is no capacity left so $e_{1,i} = e_{2,i} = \dots = e_{d,i} = 0$. If we open it, then $e_{j,i} = [u_i]_j$
 531, i.e. the j -th machine will hold exactly $[u_i]_j$ unused time slots in the interval. Now the
 532 problem becomes we want to open k special time slots such that the resulting configuration
 533 can fit all jobs from S_3 . Lemma 15 implies that it is equivalent to choosing k vectors
 534 from $(e_{1,1}, \dots, e_{d,1}) = u_1, \dots, (e_{1,n}, \dots, e_{d,n}) = u_n$ such that $\sum_{i \leq j} e_i \geq \sum_{i \leq j} [v]_i$ for every
 535 $j \in [1, d]$, which is exactly the definition of our prefix sum cover problem. Note that the
 536 ordering requirement comes from the fact we have ordering requirement in Lemma 15 and
 537 the positiveness of u comes from the fact that the machine from 1 to d has 1 unused space at
 538 time slot $[(i-1)W, (i-1)W+1]$ if we open it. Since d, W are all polynomial, the interval
 539 length and the machine number p are also polynomial.

540 ——— References ———

- 541 1 Susanne Albers. Energy-efficient algorithms. *Communications of the ACM*, 53(5):86–96, 2010.
- 542 2 Jessica Chang, Harold N. Gabow, and Samir Khuller. A model for minimizing active processor
 543 time. *Algorithmica*, 70(3):368–405, 2014.
- 544 3 Jessica Chang, Samir Khuller, and Koyel Mukherjee. LP rounding and combinatorial algorithms
 545 for minimizing active and busy time. *J. of Scheduling*, 20(6):657–680, 2017.
- 546 4 Vincent Chau and Minming Li. *Active and Busy Time Scheduling Problem: A Survey*, pages
 547 219–229. Springer International Publishing, 2020.
- 548 5 Julia Chuzhoy, Sudipto Guha, Sanjeev Khanna, and Joseph Naor. Machine minimization for
 549 scheduling jobs with interval constraints. In *Proc. of the 45th Symposium on Foundations of*
 550 *Computer Science*, pages 81–90, 2004.
- 551 6 Gruia Călinescu and Kai Wang. A new lp rounding algorithm for the active time problem.
 552 *Journal of Scheduling*, pages 1–10, 03 2021.
- 553 7 Sami Davies, Samir Khuller, and Shirley Zhang. Balancing flow time and energy consumption.
 554 In *Proceedings of the 34th ACM Symposium on Parallelism in Algorithms and Architectures*,
 555 SPAA '22, page 369–380, New York, NY, USA, 2022. Association for Computing Machinery.
 556 doi:10.1145/3490148.3538582.
- 557 8 Sandy Irani and Kirk R. Pruhs. Algorithmic problems in power management. *SIGACT News*,
 558 36(2):63–76, 2005.
- 559 9 Frederic Koehler and Samir Khuller. Busy time scheduling on a bounded number of machines
 560 (extended abstract). In *Proc. of the 15th International Symposium on Algorithms and Data*
 561 *Structures*, pages 521–532, 2017.
- 562 10 Saurabh Kumar and Samir Khuller. Brief announcement: A greedy 2 approximation for the
 563 active time problem. In *Proc. of the 30th on Symposium on Parallelism in Algorithms and*
 564 *Architectures*, page 347–349, 2018.
- 565 11 Mozhengfu Liu and Xueyan Tang. Analysis of busy-time scheduling on heterogeneous machines.
 566 In *Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures*,

- 567 SPAA '21, page 340–350, New York, NY, USA, 2021. Association for Computing Machinery.
 568 doi:10.1145/3409964.3461795.
- 569 12 Sagnik Saha and Manish Purohit. Np-completeness of the active time scheduling problem.
 570 *arXiv preprint arXiv:2112.03255*, 2021.
- 571 13 Ward Van Heddeghem, Sofie Lambert, Bart Lannoo, Didier Colle, Mario Pickavet, and Piet
 572 Demeester. Trends in worldwide ICT electricity consumption from 2007 to 2012. *Computer*
 573 *Communications*, 50:64–76, 2014.
- 574 14 Laurence A. Wolsey. An analysis of the greedy algorithm for the submodular set covering
 575 problem. *Combinatorica*, 2(4):385–393, 1982.

576 **A** NP Completeness of the prefix sum cover problem

577 **Proof.** We will reduce set cover problem to the prefix sum cover problem. Recall $[v]_i$ is the
 578 i -th index value of vector v . Consider a set cover instance, U is the universe containing d
 579 elements, S contains n sets and k is the target integer, the set cover problem is to find at
 580 most k subsets from S such that the union of those sets is the universe U . We could use a
 581 vector $u_i \in N^d$ to represent each set of S , for each index $j \in [1, d]$, if the set contains the
 582 j -th element, then we set the j -th value of u_i to be 1. Otherwise, it will be 0. the target
 583 vector $v = \mathbf{1}^d$. Now the set cover problem is to find at most k vectors $u_{i_1}, u_{i_2}, \dots, u_{i_k}$ from
 584 u_1, u_2, \dots, u_n such that for each $j \in [1, d]$, $[\sum_{i \leq k} u_{i_i}]_j \geq [v]_j$. For technique problem, we will
 585 add 0-th index to the vector and set it to be 0 for both u and v . This won't affect our
 586 solution and this index is only helps for dealing with 1-th index.

587 Next, we will transform all vectors u_1, u_2, \dots, u_n to new vectors u'_1, u'_2, \dots, u'_n . Specifically,
 588 for each vector $u_i = ([u_i]_1, [u_i]_2, \dots, [u_i]_d)$, the new vector $u'_i = ([u'_i]_1, [u'_i]_2, \dots, [u'_i]_d)$, where
 589 $[u'_i]_j = [u_i]_j - [u_i]_{j-1} + 2 + d - j$ for all $j \in [1, d]$. Notice that $[u_i]_j$ is either 0 or 1, thus $[u'_i]_j \in$
 590 $[1, d + 2]$. Now for the target vector v , we will set the new vector $v' = ([v']_1, [v']_2, \dots, [v']_d)$
 591 such that $[v']_j = [v]_j - [v]_{j-1} + 2k + k(d - j)$. Again, since $[v]_j, [v]_{j-1} \in [0, 1]$, we have
 592 $[v']_j \in [2k - 1, kd + k + 1]$. Thus, the maximum value in the new vectors is at most $kd + k + 1$,
 593 which is at most polynomial in n and fits our requirement of prefix sum cover problem.
 594 Last, for the ordering requirement, we have $[u'_i]_j - [u'_i]_{j-1} = [u_i]_{j-2} - [u_i]_{j-1} + 1 \geq 0$ and
 595 $[v']_j - [v']_{j-1} = [v]_{j-2} - [v]_{j-1} + k \geq 0$. Now we want to show the following if-and-only-if
 596 for the reduction.

597 **If part** If we have a solution $u_{i_1}, u_{i_2}, \dots, u_{i_k}$ for the set cover problem. If the solution
 598 contains less than k vectors, we could add some vectors to the solution until k vectors, this
 599 doesn't change the solution, so we could assume we have k vectors in the solution. Now,
 600 we want to show, the new vector $u'_{i_1}, u'_{i_2}, \dots, u'_{i_k}$ is a solution for the partial sum problem.
 601 From set cover problem, we know $[\sum_{i \leq k} u_{i_i}]_j \geq [v]_j$, for $j \in [0, d]$. Our target is to show
 602 $\sum_{i' \leq j} \sum_{i \leq k} [u'_{i_i}]_{i'} \geq \sum_{i' \leq j} [v']_{i'}$, for $j \in [1, d]$. Notice that

$$\begin{aligned}
 & \sum_{i' \leq j} \sum_{i \leq k} [u'_{i_i}]_{i'} - \sum_{i' \leq j} [v']_{i'} \\
 &= \sum_{i \leq k} \sum_{i' \leq j} ([u_{i_i}]_{i'} - [u_{i_i}]_{i'-1} + 2 + d - i') - \sum_{i' \leq j} ([v]_{i'} - [v]_{i'-1} + 2k + k(d - j)) \\
 603 &= \sum_{i \leq k} ([u_{i_i}]_j + 2j + \frac{(2d - 1 - j)j}{2}) - ([v]_j + 2jk + \frac{(2d - 1 - j)jk}{2}) \\
 &= \sum_{i \leq k} [u_{i_i}]_j - [v]_j = [\sum_{i \leq k} u_{i_i}]_j - [v]_j \geq 0
 \end{aligned}$$

604 Thus, the new vectors are the solution for the partial problem.

605 **Only if** If we have a solution $u'_{l_1}, u'_{l_2}, \dots, u'_{l_k}$ for the prefix sum cover problem. Again, if
606 the solution contains less than k vectors, we add some vectors to the solution. Notice that
607 all number in the vector are non-negative, thus, the new solution is still feasible. Now, based
608 on the above equation, the vector $u_{l_1}, u_{l_2}, \dots, u_{l_k}$ is a solution for the set cover problem.

609

