



Eberle, Franziska ; Megow, Nicole ; Schewior, Kevin

### Online Throughput Maximization on Unrelated Machines: Commitment is No Burden

Journal Article as: peer-reviewed accepted version (Postprint)

DOI of this document\* (secondary publication): <https://doi.org/10.26092/elib/3189>

Publication date of this document: 01/08/2024

\* for better findability or for reliable citation

#### Recommended Citation (primary publication/Version of Record) incl. DOI:

Franziska Eberle, Nicole Megow, and Kevin Schewior. 2023. Online Throughput Maximization on Unrelated Machines: Commitment is No Burden. ACM Trans. Algorithms 19, 1, Article 10 (January 2023), 25 pages. <https://doi.org/10.1145/3569582>.

Please note that the version of this document may differ from the final published version (Version of Record/primary publication) in terms of copy-editing, pagination, publication date and DOI. Please cite the version that you actually used. Before citing, you are also advised to check the publisher's website for any subsequent corrections or retractions (see also <https://retractionwatch.com/>).

© Authors | ACM 2023. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in ACM Transactions on Algorithms, <https://doi.org/10.1145/3569582>.

This document is made available with all rights reserved.

#### Take down policy

If you believe that this document or any material on this site infringes copyright, please contact [publizieren@suub.uni-bremen.de](mailto:publizieren@suub.uni-bremen.de) with full details and we will remove access to the material.

# ONLINE THROUGHPUT MAXIMIZATION ON UNRELATED MACHINES: COMMITMENT IS NO BURDEN\*

FRANZISKA EBERLE<sup>†</sup>, NICOLE MEGOW<sup>†</sup>, AND KEVIN SCHEWIOR<sup>‡</sup>

**Abstract.** We consider a fundamental online scheduling problem in which jobs with processing times and deadlines arrive online over time at their release dates. The task is to determine a feasible preemptive schedule on a single or multiple possibly unrelated machines that maximizes the number of jobs that complete before their deadline. Due to strong impossibility results for competitive analysis on a single machine, we require that jobs contain some *slack*  $\varepsilon > 0$ , which means that the feasible time window for scheduling a job is at least  $1 + \varepsilon$  times its processing time on each eligible machine. Our contribution is two-fold: (i) We give the first non-trivial online algorithms for throughput maximization on unrelated machines, and (ii), this is the main focus of our paper, we answer the question on how to handle commitment requirements which enforce that a scheduler has to guarantee at a certain point in time the completion of admitted jobs. This is very relevant, e.g., in providing cloud-computing services, and disallows last-minute rejections of critical tasks. We present an algorithm for unrelated machines that is  $\Theta(\frac{1}{\varepsilon})$ -competitive when the scheduler must commit upon starting a job. Somewhat surprisingly, this is the same optimal performance bound (up to constants) as for scheduling without commitment on a single machine. If commitment decisions must be made before a job's slack becomes less than a  $\delta$ -fraction of its size, we prove a competitive ratio of  $\mathcal{O}(\frac{1}{\varepsilon - \delta})$  for  $0 < \delta < \varepsilon$ . This result nicely interpolates between commitment upon starting a job and commitment upon arrival. For the latter commitment model, it is known that no (randomized) online algorithm admits any bounded competitive ratio. While we mainly focus on scheduling without migration, our results also hold when comparing against a migratory optimal solution in case of identical machines.

**Key words.** Deadline scheduling, throughput, online algorithms, competitive analysis, unrelated machines, migration

**AMS subject classifications.** 68W27, 90B35, 68W40, 68Q25

**1. Introduction.** We consider the following online scheduling problem: there are given  $m$  unrelated parallel machines. Jobs from an unknown job set arrive online over time at their *release dates*  $r_j$ . Each job  $j$  has a *deadline*  $d_j$  and a *processing time*  $p_{ij} \in \mathbb{R}_+ \cup \{\infty\}$ , which is the execution time of  $j$  when processing on machine  $i$ ; both job parameters become known to an algorithm at job arrival. We denote a machine  $i$  with  $p_{ij} < \infty$  as *eligible* for job  $j$ . If all machines are identical,  $p_{ij} = p_j$  holds for every job  $j$ , and we omit the index  $i$ . When scheduling these jobs or a subset of them, we allow *preemption*, i.e., the processing of a job can be interrupted at any time and may resume later without any additional cost. We mainly study scheduling without *migration* which means that a job must run completely on one machine. In case that we allow migration, a preempted job can resume processing on any machine, but no job can run simultaneously on two or more machines.

In a feasible schedule, two jobs are never processing at the same time on the same machine. A job is said to *complete* if it receives  $p_{ij}$  units of processing time

---

\*Submitted to the editors June 9, 2021. A preliminary version of the results on online scheduling with commitment on parallel identical machines was published at the European Symposium of Algorithms 2020 [15].

**Funding:** The work of the second author was partially supported by the German Science Foundation (DFG) under contract ME 3825/1. The work of the third author was partially supported by the DAAD within the PRIME program using funds of BMBF and the EU Marie Curie Actions.

<sup>†</sup>Faculty of Mathematics and Computer Science, University of Bremen, Germany ([franziska.eberle@posteo.de](mailto:franziska.eberle@posteo.de), [nicole.megow@uni-bremen.de](mailto:nicole.megow@uni-bremen.de)).

<sup>‡</sup>Department of Mathematics and Computer Science, Universität zu Köln, Cologne, Germany ([kschewior@gmail.com](mailto:kschewior@gmail.com)).

41 on machine  $i$  within the interval  $[r_j, d_j)$  if  $j$  is processed by machine  $i$ . The number  
 42 of completed jobs in a feasible schedule is called *throughput*. The task is to find a  
 43 feasible schedule with maximum throughput. We refer to this problem as *throughput*  
 44 *maximization*.

45 As jobs arrive online and scheduling decisions are irrevocable, we cannot hope to  
 46 find an optimal schedule even when scheduling on a single machine [12]. To assess  
 47 the performance of online algorithms, we resort to standard *competitive analysis*.  
 48 This means, we compare the throughput of an online algorithm with the throughput  
 49 achievable by an optimal offline algorithm that knows the job set in advance.

50 On a single machine, it is well-known that “tight” jobs with  $d_j - r_j \approx p_j$  prohibit  
 51 competitive online decision making as jobs must start immediately and do not leave  
 52 a chance for observing online arrivals [7]. Thus, it is commonly required that jobs  
 53 contain some *slack*  $\varepsilon > 0$ , i.e., every job  $j$  satisfies  $d_j - r_j \geq (1 + \varepsilon)p_j$ . In the more  
 54 general setting with unrelated machines, we assume that each job  $j$  satisfies  $d_j - r_j \geq$   
 55  $(1 + \varepsilon)p_{ij}$  for each machine  $i$  that is eligible for  $j$ , i.e., each machine  $i$  with  $p_{ij} < \infty$ .  
 56 The competitive ratio of our online algorithm will be a function of  $\varepsilon$ ; the greater the  
 57 slack, the better should the performance of our algorithm be. This slackness parameter  
 58 has been considered in a multitude of previous work, e.g., in [2, 5, 10, 17, 18, 31, 34].  
 59 Other results for scheduling with deadlines use speed scaling, which can be viewed as  
 60 another way to add slack to the schedule, see, e.g., [1, 3, 22, 24, 32].

61 In this paper, we focus on the question how to handle *commitment* requirements  
 62 in online throughput maximization. Modeling commitment addresses the issue that  
 63 a high-throughput schedule may abort jobs close to their deadlines in favor of many  
 64 shorter and more urgent tasks [16], which may not be acceptable for the job owner.  
 65 Consider a company that starts outsourcing mission-critical processes to external  
 66 clouds and that needs a guarantee that jobs complete before a certain time point when  
 67 they cannot be moved to another computing cluster anymore. In other situations, a  
 68 commitment to complete jobs might be required even earlier just before starting the  
 69 job, e.g., for a faultless copy of a database [10].

70 Different commitment models have been formalized [2, 10, 31]. The requirement to  
 71 commit at a job’s release date has been ruled out for online throughput maximization  
 72 by strong impossibility results (even for randomized algorithms) [10]. We distinguish  
 73 two commitment models.

- 74 (i) *Commitment upon job admission*: an algorithm may discard a job any time  
 75 before its start, we say its admission. This reflects a situation such as the  
 76 faultless copy of a database.
- 77 (ii)  $\delta$ -*commitment*: given  $0 < \delta < \varepsilon$ , an algorithm must commit to complete  
 78 a job while the job’s remaining slack is at least a  $\delta$ -fraction of its original  
 79 processing time. This models an early enough commitment (parameterized  
 80 by  $\delta$ ) for mission-critical jobs. For identical parallel machines, the latest time  
 81 for committing to job  $j$  is then  $d_j - (1 + \delta)p_j$ . When given unrelated machines,  
 82 such a commitment model might be arguably less relevant. We consider it  
 83 only for non-migratory schedules and include also the choice of a processor  
 84 in the commitment; we define the latest time point for committing to job  $j$   
 85 as  $d_j - (1 + \delta)p_{ij}$  when processing  $j$  on machine  $i$ .

86 Recently, a first unified approach has been presented for these models for a single  
 87 machine [10]. In this and other works [2, 31], there remained gaps in the performance  
 88 bounds and it was left open whether scheduling with commitment is even “harder”  
 89 than without commitment. Moreover, it remained unsettled whether the problem is  
 90 tractable on multiple identical or even heterogeneous machines.

91 In this work, we give tight results for online throughput maximization on un-  
 92 related parallel machines and answer the “hardness” question to the negative. We  
 93 give an algorithm that achieves the provably best competitive ratio (up to constant  
 94 factors) for the aforementioned commitment models. Somewhat surprisingly, we show  
 95 that the same competitive ratio of  $\mathcal{O}(\frac{1}{\varepsilon})$  can be achieved for both, scheduling *without*  
 96 commitment and *with* commitment upon admission. For unrelated machines, this  
 97 is the first nontrivial result for online throughput maximization with and without  
 98 commitment. For identical parallel machines, this is the first online algorithm with  
 99 bounded competitive ratio for arbitrary slack parameter  $\varepsilon$ . Interestingly, for this  
 100 machine environment, our algorithm does not require job migration in order to be  
 101 competitive against a migratory algorithm.

102 **1.1. Related work.** Preemptive online scheduling and admission control have  
 103 been studied rigorously. There are several results regarding the impact of deadlines  
 104 on online scheduling; see, e.g., [6, 17, 18] and references therein. In the following we  
 105 give an overview of the literature focused on (online) throughput maximization.

106 *Offline scheduling.* In case that the jobs and their characteristics are known  
 107 to the scheduler in advance, the notion of commitment is irrelevant as an offline  
 108 algorithm only starts jobs that will be completed on time; there is no benefit in  
 109 starting jobs without completing them. The offline problem is well understood: For  
 110 throughput maximization on a single machine, there is a polynomial-time algorithm  
 111 by Lawler [29]. The model where jobs have *weights* and the task is to maximize the  
 112 total weight of jobs completed on time (*weighted throughput*) is NP-hard and we do  
 113 not expect polynomial time algorithms. The algorithm by Lawler solves this problem  
 114 optimally in time  $\mathcal{O}(n^5 w_{\max})$ , where  $w_{\max} = \max_j w_j$ , and can be used to design a  
 115 fully polynomial-time approximation scheme (FPTAS) [33].

116 When given multiple identical machines, (unweighted) throughput maximiza-  
 117 tion becomes NP-hard even for identical release dates [30]. Kalyanasundaram and  
 118 Pruhs [25] show a 6-approximate reduction to the single-machine problem which im-  
 119 plies a  $(6 + \delta)$ -approximation algorithm for weighted throughput maximization on  
 120 identical parallel machines, for any  $\delta > 0$ , using the FPTAS for the single-machine  
 121 problem [33]. Preemptive throughput maximization on unrelated machines is much  
 122 less understood from an approximation point of view. The problem is known to be  
 123 strongly NP-hard [14], even without release dates [35]. We are not aware of any ap-  
 124 proximation results for preemptive throughput maximization on unrelated machines.  
 125 The situation is different for non-preemptive scheduling. In this case, throughput  
 126 maximization is MAX-SNP hard [4] and several approximation algorithms for this  
 127 general problem as well as for identical parallel machines and other special cases are  
 128 known; see, e.g., [4, 9, 21].

129 *Online scheduling without commitment.* For single-machine throughput maxi-  
 130 mization, Baruah, Haritsa, and Sharma [6] show that, in general, no deterministic  
 131 online algorithm achieves a bounded competitive ratio. Thus, their result justifies our  
 132 assumption on  $\varepsilon$ -slackness of each job. Moreover, they consider special cases such as  
 133 unit-size jobs or agreeable deadlines where they provide constant-competitive algo-  
 134 rithms even without further assumptions on the slack of the jobs. Here, deadlines are  
 135 agreeable if  $r_j \leq r_{j'}$  for two jobs  $j$  and  $j'$  implies  $d_j \leq d_{j'}$ . In our prior work [10], we  
 136 develop a  $\Theta(\frac{1}{\varepsilon})$ -competitive algorithm and show a matching lower bound for deter-  
 137 ministic algorithms. While this is ruled out for deterministic algorithms, Kalyanasun-  
 138 daram and Pruhs [26] give a *randomized*  $\mathcal{O}(1)$ -competitive algorithm for throughput  
 139 maximization on a single machine without slackness assumption.

140 For maximizing weighted throughput, Lucier et al. [31] give an  $\mathcal{O}(\frac{1}{\varepsilon^2})$ -competitive  
 141 online algorithm for scheduling on identical parallel machines. In a special case of this  
 142 problem, called *machine utilization* the goal is to maximize the total processing time  
 143 of completed jobs. This problem is much more tractable. On a single machine, Baruah  
 144 et al. [7, 8] provide a best-possible online algorithm achieving a competitive ratio of 4,  
 145 even without any slackness assumptions. Baruah and Haritsa [5] are the first to inves-  
 146 tigate the problem under the assumption of  $\varepsilon$ -slack and give a  $\frac{1+\varepsilon}{\varepsilon}$ -competitive algo-  
 147 rithm which is asymptotically best possible. For parallel identical machines (though  
 148 without migration), DasGupta and Palis [11] give a simple greedy algorithm that  
 149 achieves the same performance guarantee of  $\frac{1+\varepsilon}{\varepsilon}$  and give an asymptotically match-  
 150 ing lower bound. Schwiegelshohn and Schwiegelshohn [34] show that migration helps  
 151 an online algorithm and improves the competitive ratio to  $\mathcal{O}(\sqrt[m]{1/\varepsilon})$  for  $m$  machines.

152 In a line of research without slackness assumption, Baruah et al. [8] show a lower  
 153 bound of  $(1 + \sqrt{k})^2$  for deterministic single-machine algorithms, where  $k = \frac{\max_j w_j/p_j}{\min_j w_j/p_j}$   
 154 is the *importance ratio* of a given instance. Koren and Shasha give a matching upper  
 155 bound [28] and generalize it to  $\Theta(\ln k)$  for parallel machines if  $k > 1$  [27].

156 *Online scheduling with commitment upon job arrival.* In our prior work [10], we  
 157 rule out bounded competitive ratios for any (even randomized) online algorithm for  
 158 throughput maximization with commitment upon job arrival, even on a single ma-  
 159 chine. Previously, such impossibility results were only shown exploiting weights [31].

160 Again, the special case  $w_j = p_j$ , or machine utilization, is much more tractable  
 161 than weighted or unweighted throughput maximization. A simple greedy algorithm  
 162 already achieves the best possible competitive ratio  $\frac{1+\varepsilon}{\varepsilon}$  on a single machine, even for  
 163 commitment upon arrival, as shown by DasGupta and Palis [11] and the matching  
 164 lower bound by Garay et al. [17]. For scheduling with commitment upon arrival on  $m$   
 165 parallel identical machines, there is an  $\mathcal{O}(\sqrt[m]{1/\varepsilon})$ -competitive algorithm and an al-  
 166 most matching lower bound by Schwiegelshohn and Schwiegelshohn [34]. Surprisingly,  
 167 this model also allows for bounded competitive ratios when preemption is not allowed.  
 168 In this setting, Goldwasser and Kerbikov [19] give a best possible  $(2 + \frac{1}{\varepsilon})$ -competitive  
 169 algorithm on a single machine. Very recently, Jamalabadi, Schwiegelshohn, and  
 170 Schwiegelshohn [23] extend this model to parallel machines; their algorithm is near  
 171 optimal with a performance guarantee approaching  $\ln \frac{1}{\varepsilon}$  as  $m$  tends to infinity.

172 *Online scheduling with commitment upon admission and  $\delta$ -commitment.* In our  
 173 previous work [10], we design an online single-machine algorithm, called the *region*  
 174 *algorithm*, that simultaneously (with the respective choice of parameters) achieves the  
 175 first non-trivial upper bounds for both commitment models. For commitment upon  
 176 job admission, our bound is  $\mathcal{O}(\frac{1}{\varepsilon^2})$ , and in the  $\delta$ -commitment model it is  $\mathcal{O}(\frac{\varepsilon}{(\varepsilon-\delta)\delta^2})$ ,  
 177 for  $0 < \delta < \varepsilon$ . For scheduling on identical parallel machines and commitment upon  
 178 admission, Lucier et al. [31] give a heuristic that empirically performs very well but  
 179 for which they cannot show a rigorous worst-case bound. In fact, Azar et al. [2] show  
 180 that no bounded competitive ratio is possible for weighted throughput maximization  
 181 for small  $\varepsilon$ . For  $\delta = \frac{\varepsilon}{2}$  in the  $\delta$ -commitment model, they design (in the context  
 182 of truthful mechanisms) an online algorithm for weighted throughput maximization  
 183 that is  $\Theta(\frac{1}{\sqrt[3]{1+\varepsilon}-1} + \frac{1}{(\sqrt[3]{1+\varepsilon}-1)^2})$ -competitive if the slack  $\varepsilon$  is sufficiently large, i.e.,  
 184 if  $\varepsilon > 3$ . For weighted throughput, this condition on the slack is necessary as is shown  
 185 by a strong general lower bound, even on a single machine [10]. For the unweighted  
 186 setting, we give the first rigorous upper bound for arbitrary  $\varepsilon$  in this paper for both  
 187 models, commitment upon admission and  $\delta$ -commitment, in the identical and even in  
 188 the unrelated machine environment.

189 Machine utilization is again better understood. As commitment upon arrival is  
 190 more restrictive than commitment upon admission and  $\delta$ -commitment, the previously  
 191 mentioned results immediately carry over and provide bounded competitive ratios.

192 **1.2. Our results and techniques.** Our main result is an algorithm that com-  
 193 putes a non-migratory schedule that is best possible (up to constant factors) for  
 194 online throughput maximization with and without commitment on identical parallel  
 195 machines and, more generally, on unrelated machines. This is the first non-trivial  
 196 online result for unrelated machines and it closes gaps for identical parallel machines.  
 197 Our algorithm is universally applicable (by setting parameters properly) to both com-  
 198 mitment models as well as scheduling without commitment.

199 **THEOREM 1.1.** *Consider throughput maximization on unrelated machines with-*  
 200 *out migration. There is an  $\mathcal{O}(\frac{1}{\varepsilon-\delta'})$ -competitive non-migratory online algorithm for*  
 201 *scheduling with commitment, where  $\delta' = \frac{\varepsilon}{2}$  in the model with commitment upon ad-*  
 202 *mission and  $\delta' = \max\{\delta, \frac{\varepsilon}{2}\}$  in the  $\delta$ -commitment model.*

203 For scheduling with commitment upon admission, this is (up to constant factors)  
 204 an optimal online algorithm with competitive ratio  $\Theta(\frac{1}{\varepsilon})$ , matching the lower bound  
 205 of  $\Omega(\frac{1}{\varepsilon})$  for  $m = 1$  [10]. For scheduling with  $\delta$ -commitment, our result interpolates  
 206 between the models with commitment upon starting a job and commitment upon  
 207 arrival. If  $\delta \leq \frac{\varepsilon}{2}$ , the competitive ratio is  $\Theta(\frac{1}{\varepsilon})$ , which is again best possible [10].  
 208 For  $\delta \rightarrow \varepsilon$ , the commitment requirements essentially implies commitment upon job  
 209 arrival which has unbounded competitive ratio [10].

210 In our analysis, we compare a non-migratory schedule, obtained by our algorithm,  
 211 with an optimal non-migratory schedule. However, in the case of identical machines  
 212 the throughput of an optimal migratory schedule can only be larger by a constant  
 213 factor than the throughput of an optimal non-migratory schedule. In fact, Kalyana-  
 214 sundaram and Pruhs [25] showed that this factor is at most  $\frac{6m-5}{m}$ . Thus, the com-  
 215 petitive ratio for our non-migratory algorithm, when applied to identical machines,  
 216 holds (up to this constant factor) also in a migratory setting.

217 **COROLLARY 1.2.** *Consider throughput maximization with or without migration on*  
 218 *parallel identical machines. There is an  $\mathcal{O}(\frac{1}{\varepsilon-\delta'})$ -competitive non-migratory online al-*  
 219 *gorithm for scheduling with commitment, where  $\delta' = \frac{\varepsilon}{2}$  in the model with commitment*  
 220 *upon admission and  $\delta' = \max\{\delta, \frac{\varepsilon}{2}\}$  in the  $\delta$ -commitment model.*

221 The challenge in online scheduling with commitment is that, once we committed  
 222 to complete a job, the remaining slack of this job has to be spent very carefully.  
 223 The key component is a job admission scheme which is implemented by different  
 224 parameters. The high-level objectives are:

- 225 (i) never start a job for the first time if its remaining slack is too small (param-  
 226 eter  $\delta$ ),
- 227 (ii) during the processing of a job, admit only significantly shorter jobs (param-  
 228 eter  $\gamma$ ), and
- 229 (iii) for each admitted shorter job, block some time period (parameter  $\beta$ ) during  
 230 which no other jobs of similar size are accepted.

231 While the first two goals are quite natural and have been used before in the single  
 232 and identical machine setting [10,31], the third goal is crucial for our new tight result.  
 233 The intuition is the following: Think of a single eligible machine in a non-migratory  
 234 schedule. Suppose we committed to complete a job with processing time 1 and have  
 235 only a slack of  $\mathcal{O}(\varepsilon)$  left before the deadline of this job. Suppose that  $c$  substantially  
 236 smaller jobs of size  $\frac{1}{c}$  arrive where  $c$  is the competitive ratio we aim for. On the one

237 hand, if we do not accept any of them, we cannot hope to achieve  $c$ -competitiveness.  
 238 On the other hand, accepting too many of them fills up the slack and, thus, leaves no  
 239 room for even smaller jobs. The idea is to keep the flexibility for future small jobs by  
 240 only accepting an  $\varepsilon$ -fraction of jobs of similar size (within a factor two).

241 We distinguish two time periods that guide the acceptance decisions. During  
 242 the *scheduling interval* of a job  $j$ , we have a more restrictive acceptance scheme that  
 243 ensures the completion of  $j$  whereas in the *blocking period* we guarantee the comple-  
 244 tion of previously accepted jobs. We call our algorithm *blocking* algorithm. This  
 245 acceptance scheme is much more refined than the one of the known region algorithm  
 246 in [10] that uses one long region with a uniform acceptance threshold and is then too  
 247 conservative in accepting jobs.

248 Given that we consider the non-migratory version of the problem, a generalization  
 249 from a single to multiple machines seems natural. It is interesting, however, that such  
 250 a generalization works, essentially on a per-machine basis, even for unrelated machines  
 251 and comes at no loss in the competitive ratio.

252 Clearly, scheduling with commitment is more restrictive than without commit-  
 253 ment. Therefore, our algorithm is also  $O(\frac{1}{\varepsilon})$ -competitive for maximizing the through-  
 254 put on unrelated machines without any commitment requirements. Again, this is  
 255 optimal (up to constant factors) as it matches the lower bound on the competitive  
 256 ratio for deterministic online algorithms on a single machine [10].

257 **COROLLARY 1.3.** *There is a  $\Theta(\frac{1}{\varepsilon})$ -competitive algorithm for online throughput*  
 258 *maximization on unrelated machines without commitment requirements and without*  
 259 *migration.*

260 However, for scheduling without commitment, we are able to generalize the sim-  
 261 pler region algorithm presented for the single-machine problem in [10] to scheduling  
 262 on unrelated machines.

263 **THEOREM 1.4.** *A generalization of the region algorithm is  $\Theta(\frac{1}{\varepsilon})$ -competitive for*  
 264 *online throughput maximization on unrelated machines without commitment require-*  
 265 *ments and without migration.*

266 Besides presenting a simpler algorithm for throughput maximization without com-  
 267 mitment, we show this result to present an additional application of our technical  
 268 findings for the analysis of the blocking algorithm. We give details later. On a high  
 269 level, we show a key lemma on the size of non-admitted jobs for a big class of online  
 270 algorithms which results in an upper bound on the throughput of an optimal (offline)  
 271 non-migratory algorithm. This key lemma can be used in the analysis of both algo-  
 272 rithms, blocking and region. In fact, also the analysis of the original region algorithm  
 273 for a single machine [10] becomes substantially easier.

274 In case of identical machines, again, we can apply the result by Kalyanasundaram  
 275 and Pruhs [25] that states that the throughput of an optimal migratory schedule is  
 276 larger by at most a constant factor than the throughput of an optimal non-migratory  
 277 schedule. Thus, the result in [Theorem 1.4](#) holds also in a migratory setting when  
 278 scheduling on identical machines.

279 **COROLLARY 1.5.** *A generalization of the region algorithm is  $\Theta(\frac{1}{\varepsilon})$ -competitive for*  
 280 *online throughput maximization on multiple identical machines without commitment*  
 281 *requirements, with and without migration.*

282 **Outline of the paper.** In [Section 2](#), we describe and outline the analysis of our  
 283 new non-migratory algorithm. It consists of two parts, which are detailed in [Sections 3](#)

284 and 4: firstly, we argue that all jobs admitted by our algorithm can complete by their  
 285 deadline and, secondly, we prove that we admit “sufficiently many” jobs. In Section 5,  
 286 we generalize the known region algorithm, developed for a single machine in our prior  
 287 work [10], to a non-migratory algorithm without commitment on unrelated machines.  
 288 We show how to apply a new key technique developed for the analysis in Section 4  
 289 to analyze it and prove the same competitive ratio (up to constant factors) as for a  
 290 single machine.

291 **2. The blocking algorithm.** In this section, we describe the *blocking algorithm*  
 292 for scheduling with commitment. We assume that the slackness constant  $\varepsilon > 0$  and,  
 293 in the  $\delta$ -commitment model,  $\delta \in (0, \varepsilon)$  are given. If  $\delta$  is not part of the input or  
 294 if  $\delta \leq \frac{\varepsilon}{2}$ , then we set  $\delta = \frac{\varepsilon}{2}$ .

295 The algorithm never migrates jobs between machines, i.e., a job is only processed  
 296 by the machine that initially started to process it. In this case, we say the job has been  
 297 *admitted* to this machine. Moreover, our algorithm commits to completing a job upon  
 298 admission (even in the  $\delta$ -commitment model). Hence, its remaining slack has to be  
 299 spent very carefully on admitting other jobs to still be competitive. As our algorithm  
 300 does not migrate jobs, it transfers the admission decision to the shortest admitted and  
 301 not yet completed job on each machine. A job only admits significantly shorter jobs  
 302 and prevents the admission of too many jobs of similar size. To this end, the algorithm  
 303 maintains two types of intervals for each admitted job, a *scheduling interval* and a  
 304 *blocking period*. A job can only be processed in its scheduling interval. Thus, it has  
 305 to complete in this interval while admitting other jobs. Job  $j$  only admits jobs that  
 306 are smaller by a factor of at least  $\gamma = \frac{\delta}{16} < 1$ . For an admitted job  $k$ , job  $j$  creates  
 307 a blocking period of length at most  $\beta p_{ik}$ , where  $\beta = \frac{16}{\delta}$ , which blocks the admission  
 308 of similar-length jobs (cf. Figure 1). The scheduling intervals and blocking periods of  
 309 jobs admitted by  $j$  will always be pairwise disjoint and completely contained in the  
 310 scheduling interval of  $j$ .




FIG. 1. *Scheduling interval, blocking period, and processing intervals*

311 *Scheduling jobs.* Independent of the admission scheme, the blocking algorithm  
 312 follows the SHORTEST PROCESSING TIME (SPT) order for the set of uncompleted  
 313 jobs assigned to a machine. SPT ensures that a job  $j$  has highest priority in the  
 314 blocking periods of any job  $k$  admitted by  $j$ .

315 *Admitting jobs.* The algorithm keeps track of *available* jobs at any time point  $\tau$ .  
 316 A job  $j$  with  $r_j \leq \tau$  is called available for machine  $i$  if it has not yet been admitted to  
 317 a machine by the algorithm and its deadline is not too close, i.e.,  $d_j - \tau \geq (1 + \delta)p_{ij}$ .

318 Whenever a job  $j$  is available for machine  $i$  at a time  $\tau$  such that time  $\tau$  is not  
 319 contained in the scheduling interval of any other job admitted to  $i$ , the shortest such  
 320 job  $j$  is immediately admitted to machine  $i$  at time  $a_j := \tau$ , creating the scheduling  
 321 interval  $S(j) = [a_j, e_j)$ , where  $e_j = a_j + (1 + \delta)p_{ij}$  and an empty blocking period  $B(j) =$   
 322  $\emptyset$ . In general, however, the blocking period of a job  $j$  is a finite union of time intervals



323 associated with  $j$ , and its size is the sum of lengths of the intervals, denoted by  $|B(j)|$ .  
 324 Both, blocking period and scheduling interval, depend on machine  $i$  but we omit  $i$   
 325 from the notation as it is clear from the context; both periods are created after job  $j$   
 326 has been assigned to machine  $i$ .

327 Four types of events trigger a decision of the algorithm at time  $\tau$ : the release of a  
 328 job, the end of a blocking period, the end of a scheduling interval, and the admission  
 329 of a job. In any of these four cases, the algorithm calls the *admission routine*. This  
 330 subroutine iterates over all machines  $i$  and checks if  $j$ , the shortest job on  $i$  whose  
 331 scheduling interval contains  $\tau$ , can admit the currently shortest job  $j^*$  available for  
 332 machine  $i$ .

333 To this end, any admitted job  $j$  checks whether  $p_{ij^*} < \gamma p_{ij}$ . Only such jobs qualify  
 334 for admission by  $j$ . Upon admission by  $j$ , job  $j^*$  obtains two disjoint consecutive  
 335 intervals, the *scheduling interval*  $S(j^*) = [a_{j^*}, e_{j^*})$  and the *blocking period*  $B(j^*)$  of  
 336 size at most  $\beta p_{ij^*}$ . At the admission of job  $j^*$ , the blocking period  $B(j^*)$  is planned  
 337 to start at  $e_{j^*}$ , the end of  $j^*$ 's scheduling interval. During  $B(j^*)$ , job  $j$  only admits  
 338 jobs  $k$  with  $p_{ik} < \frac{1}{2} p_{ij^*}$ .

339 Hence, when job  $j$  decides if it admits the currently shortest available job  $j^*$  at  
 340 time  $\tau$ , it makes sure that  $j^*$  is sufficiently small and that no job  $k$  of similar (or  
 341 even smaller) processing time is blocking  $\tau$ , i.e., it verifies that  $\tau \notin B(k)$  for all jobs  $k$   
 342 with  $p_{ik} \leq 2p_{ij^*}$  admitted to the same machine. In this case, we say that  $j^*$  is a *child*  
 343 of  $j$  and that  $j$  is the *parent* of  $j^*$ , denoted by  $\pi(j^*) = j$ . If job  $j^*$  is admitted at  
 344 time  $\tau$  by job  $j$ , the algorithm sets  $a_{j^*} = \tau$  and  $e_{j^*} = a_{j^*} + (1 + \delta)p_{ij^*}$  and assigns  
 345 the scheduling interval  $S(j^*) = [a_{j^*}, e_{j^*})$  to  $j^*$ .

346 If  $e_{j^*} \leq e_j$ , the routine sets  $f_{j^*} = \min\{e_j, e_{j^*} + \beta p_{ij^*}\}$  which determines  $B(j^*) =$   
 347  $[e_{j^*}, f_{j^*})$ . As the scheduling and blocking periods of children  $k$  of  $j$  are supposed to  
 348 be disjoint, we have to *update the blocking periods*. First consider the job  $k$  with  $p_{ik} >$   
 349  $2p_{ij^*}$  admitted to the same machine whose blocking period contains  $\tau$  (if it exists), and  
 350 let  $[e'_k, f'_k)$  be the maximal interval of  $B(k)$  containing  $\tau$ . We set  $f''_k = \min\{e_j, f'_k +$   
 351  $(1 + \delta + \beta)p_{ij^*}\}$  and replace the interval  $[e'_k, f'_k)$  by  $[e'_k, \tau) \cup [\tau + (1 + \delta + \beta)p_{ij^*}, f''_k)$ . For  
 352 all other jobs  $k$  with  $B(k) \cap [\tau, \infty) \neq \emptyset$  admitted to the same machine, we replace the  
 353 remaining part of their blocking period  $[e'_k, f'_k)$  by  $[e'_k + (1 + \delta + \beta)p_{ij^*}, f''_k)$  where  $f''_k :=$   
 354  $\min\{e_j, f'_k + (1 + \delta + \beta)p_{ij^*}\}$ . In this update, we follow the convention that  $[e, f) = \emptyset$   
 355 if  $f \leq e$ . Observe that the length of the blocking period might decrease due to such  
 356 updates.

357 Note that  $e_{j^*} > e_j$  is also possible as  $j$  does not take the end of its own scheduling  
 358 interval  $e_j$  into account when admitting jobs. Thus, the scheduling interval of  $j^*$   
 359 would end outside the scheduling interval of  $j$  and inside the blocking period of  $j$ .  
 360 During  $B(j)$ , the parent  $\pi(j)$  of  $j$ , did not allocate the interval  $[e_j, e_{j^*})$  for completing  
 361 jobs admitted by  $j$  but for ensuring its own completion. Hence, the completion of  
 362 both  $j^*$  and  $\pi(j)$  is not necessarily guaranteed anymore. To prevent this, we *modify*  
 363 *all scheduling intervals*  $S(k)$  (including  $S(j)$ ) that contain time  $\tau$  of jobs admitted to  
 364 the same machine as  $j^*$  and their blocking periods  $B(k)$ . For each job  $k$  admitted to  
 365 the same machine with  $\tau \in S(k)$  (including  $j$ ) and  $e_{j^*} > e_k$ , we set  $e_k = e_{j^*}$ . We also  
 366 update their blocking periods (in fact, single intervals) to reflect their new starting  
 367 points. If the parent  $\pi(k)$  of  $k$  does not exist,  $B(k)$  remains empty; otherwise we  
 368 set  $B(k) := [e_k, f_k)$  where  $f_k = \min\{e_{\pi(k)}, e_k + \beta p_{ik}\}$ . Note that, after this update,  
 369 the blocking periods of any but the largest such job will be empty. Moreover, the just  
 370 admitted job  $j^*$  does not get a blocking period in this special case.

371 During the analysis of the algorithm, we show that any admitted job  $j$  still com-  
 372 pletes before  $a_j + (1 + \delta)p_{ij}$  and that  $e_j \leq a_j + (1 + 2\delta)p_{ij}$  holds in retrospect for all

373 admitted jobs  $j$ . Thus, any job  $j$  that admits another job  $j^*$  tentatively assigns this  
 374 job a scheduling interval of length  $(1+\delta)p_{ij^*}$  but, for ensuring its own completion, it is  
 375 prepared to lose  $(1+2\delta)p_{ij^*}$  time units of its scheduling interval  $S(j)$ . We summarize  
 376 the blocking algorithm in the following.

---

**Algorithm** Blocking algorithm
 

---

**Scheduling Routine:** At all times  $\tau$  and on all machines  $i$ , run the job with shortest processing time that has been admitted to  $i$  and has not yet completed.

**Event:** Release of a new job at time  $\tau$   
 Call Admission Routine.

**Event:** End of a blocking period or scheduling interval at time  $\tau$   
 Call Admission Routine.

**Admission Routine:**
 $i \leftarrow 1$ 
 $j^* \leftarrow$  a shortest job available at  $\tau$  for machine  $i$ , i.e.,  $j^* \in \arg \min\{p_{ij} \mid j \in \mathcal{J}, r_j \leq \tau \text{ and } d_j - \tau \geq (1+\delta)p_{ij}\}$ 
**while**  $i \leq m$  **do**
 $K \leftarrow$  the set of jobs on machine  $i$  whose scheduling intervals contain  $\tau$ 
**if**  $K = \emptyset$  **then**

 admit job  $j^*$  to machine  $i$ ,  $a_{j^*} \leftarrow \tau$ , and  $e_{j^*} \leftarrow a_{j^*} + (1+\delta)p_{ij^*}$ 
 $S(j^*) \leftarrow [a_{j^*}, e_{j^*})$  and  $B(j^*) \leftarrow \emptyset$ 

call Admission Routine

**else**
 $j \leftarrow \arg \min\{p_{ik} \mid k \in K\}$ 
**if**  $j^* < \gamma p_{ij}$  **and**  $\tau \notin B(j')$  for all  $j'$  admitted to  $i$  with  $p_{ij'} \leq 2p_{ij^*}$  **then**

 admit job  $j^*$  to machine  $i$ ,  $a_{j^*} \leftarrow \tau$ , and  $e_{j^*} \leftarrow a_{j^*} + (1+\delta)p_{ij^*}$ 
**if**  $e_{j^*} \leq e_j$  **then**
 $f_{j^*} \leftarrow \min\{e_j, e_{j^*} + \beta p_{ij^*}\}$ 
 $S(j^*) \leftarrow [a_{j^*}, e_{j^*})$  and  $B(j^*) \leftarrow [e_{j^*}, f_{j^*})$ 
**else**
 $S(j^*) \leftarrow [a_{j^*}, e_{j^*})$  and  $B(j^*) \leftarrow \emptyset$ 

 modify  $S(k)$  and  $B(k)$  for  $k \in K$ 

 update  $B(j')$  for  $j'$  admitted to machine  $i$  with  $B(j') \cap [\tau, \infty) \neq \emptyset$ 

call Admission Routine

**end if**
**else**
 $i \leftarrow i + 1$ 
 $j^* \leftarrow$  a shortest job available at  $\tau$  for machine  $i$ , i.e.,  $j^* \in \arg \min\{p_{ij} \mid j \in$ 
 $\mathcal{J}, r_j \leq \tau \text{ and } d_j - \tau \geq (1+\delta)p_{ij}\}$ 
**end if**
**end if**
**end while**


---

377 **Roadmap for the analysis.** During the analysis, it is sufficient to concentrate  
 378 on instances with small slack, as also noted in [10]. For  $\varepsilon > 1$  we run the blocking  
 379 algorithm with  $\varepsilon = 1$ , which only tightens the commitment requirement, and obtain  
 380 constant competitive ratios. Thus, we assume  $0 < \varepsilon \leq 1$ . For  $0 < \delta < \varepsilon$ , in  
 381 the  $\delta$ -commitment model an online scheduler needs to commit to the completion of a

382 job  $j$  no later than  $d_j - (1 + \delta)p_{ij}$ . Hence, committing to the completion of a job  $j$   
 383 at an earlier point in time clearly satisfies committing at a remaining slack of  $\delta p_{ij}$ .  
 384 Therefore, we may assume  $\delta \in [\frac{\varepsilon}{2}, \varepsilon)$ .

385 The blocking algorithm does not migrate any job. In the analysis, we compare  
 386 the throughput of our algorithm to the solution of an optimal non-migratory schedule.  
 387 To do so, we rely on a key design principle of the blocking algorithm, which is that,  
 388 whenever the job set admitted to a machine is fixed, the scheduling of the jobs follows  
 389 the simple SPT order. This enables us to split the analysis into two parts.

390 In the first part, we argue that the scheduling routine can handle the admitted  
 391 jobs sufficiently well. That is, every admitted jobs is completed on time; see [Section 3](#).  
 392 Here, we use that the blocking algorithm is non-migratory and consider each machine  
 393 individually.

394 For the second part, we observe that the potential admission of a new job  $j^*$   
 395 to machine  $i$  is solely based on its availability and on its size relative to  $j_i$ , the job  
 396 currently processed by machine  $i$ . More precisely, given the availability of  $j^*$ , if  $p_{ij^*} <$   
 397  $\gamma p_{ij_i}$ , the time does not belong to the blocking period of a job  $k_i$  admitted to machine  $i$   
 398 with  $p_{ij^*} \geq \frac{1}{2}p_{ik_i}$  and  $i$  is the first machine (according to machine indices) with this  
 399 property, then  $j^*$  is admitted to machine  $i$ . This implies that  $\min\{\gamma p_{ij_i}, \frac{1}{2}p_{ik_i}\}$  acts  
 400 as a *threshold*, and only available jobs with processing time less than this threshold  
 401 qualify for admission by the blocking algorithm on machine  $i$ . Hence, any available  
 402 job that the blocking algorithm does not admit has to exceed the threshold.

403 Based on this observation, we develop a general charging scheme for *any* non-  
 404 migratory online algorithm satisfying the property that, at any time  $\tau$ , the algorithm  
 405 maintains a time-dependent threshold and the shortest available job smaller than this  
 406 threshold is admitted by the algorithm. We formalize this description and analyze the  
 407 competitive ratio of such algorithms in [Section 4](#) before applying this general result  
 408 to our particular algorithm.

409 **3. Completing all admitted jobs on time.** We show that the blocking algo-  
 410 rithm finishes every admitted job on time in [Theorem 3.1](#).

411 **THEOREM 3.1.** *Let  $0 < \delta < \varepsilon$  be fixed. If  $0 < \gamma < 1$  and  $\beta \geq 1$  satisfy*

$$412 \quad (3.1) \quad \frac{\beta/2}{\beta/2 + (1 + 2\delta)} (1 + \delta - 2(1 + 2\delta)\gamma) \geq 1,$$

413 *then the blocking algorithm completes any job  $j$  admitted at  $a_j \leq d_j - (1 + \delta)p_{ij}$  on*  
 414 *time.*

415 Recall that we chose  $\gamma = \frac{\delta}{16}$  and  $\beta = \frac{16}{\delta}$ , which guarantees that Equation (3.1) is  
 416 satisfied.

417 As the blocking algorithm does not migrate jobs, it suffices to consider each ma-  
 418 chine individually in this section. The proof relies on the following observations: (i)  
 419 The sizes of jobs admitted by job  $j$  that interrupt each others' blocking periods are  
 420 geometrically decreasing, (ii) the scheduling intervals of jobs are completely contained  
 421 in the scheduling intervals of their parents, and (iii) scheduling in SPT order guaran-  
 422 tees that job  $j$  has highest priority in the blocking periods of its children. We start by  
 423 proving the following technical lemma about the length of the final scheduling interval  
 424 of an admitted job  $j$ , denoted by  $|S(j)|$ . In the proof, we use that  $\pi(k) = j$  for two  
 425 jobs  $j$  and  $k$  implies that  $p_{ik} < \gamma p_{ij}$ .

426 **LEMMA 3.2.** *Let  $0 < \delta < \varepsilon$  be fixed. If  $\gamma > 0$  satisfies  $(1 + 2\delta)\gamma \leq \delta$ , then  $|S(j)| \leq$   
 427  $(1 + 2\delta)p_{ij}$ . Moreover,  $S(j)$  contains the scheduling intervals and blocking periods of*

428 *all descendants of  $j$ .*

429 *Proof.* Consider a machine  $i$  and let  $j$  be a job admitted to machine  $i$ . By defini-  
 430 tion of the blocking algorithm, the end point  $e_j$  of the scheduling interval of job  $j$  is  
 431 only modified when  $j$  or one of  $j$ 's descendants admits another job. Let us consider  
 432 such a case: If job  $j$  admits a job  $k$  whose scheduling interval does not fit into the  
 433 scheduling interval of  $j$ , we set  $e_j = e_k = a_k + (1 + \delta)p_{ik}$  to accommodate the schedul-  
 434 ing interval  $S(k)$  within  $S(j)$ . The same modification is applied to any ancestor  $j'$  of  $j$   
 435 with  $e_{j'} < e_k$ . This implies that, after such a modification of the scheduling interval,  
 436 neither  $j$  nor any affected ancestor  $j'$  of  $j$  are the smallest jobs in their scheduling  
 437 intervals anymore. In particular, no job whose scheduling interval was modified in  
 438 such a case at time  $\tau$  is able to admit jobs after  $\tau$ . Hence, any job  $j$  can only admit  
 439 other jobs within the interval  $[a_j, a_j + (1 + \delta)p_{ij}]$ . That is,  $a_k \leq a_j + (1 + \delta)p_{ij}$  for  
 440 every job  $k$  with  $\pi(k) = j$ .

441 Thus, by induction, it is sufficient to show that  $a_k + (1 + 2\delta)p_{ik} \leq a_j + (1 + 2\delta)p_{ij}$   
 442 for admitted jobs  $k$  and  $j$  with  $\pi(k) = j$ . Note that  $\pi(k) = j$  implies  $p_{ik} < \gamma p_{ij}$ .  
 443 Hence,

$$444 \quad a_k + (1 + 2\delta)p_{ik} \leq (a_j + (1 + \delta)p_{ij}) + (1 + 2\delta)\gamma p_{ij} \leq a_j + (1 + 2\delta)p_{ij},$$

445 where the last inequality follows from the assumption  $(1 + 2\delta)\gamma \leq \delta$ . Due to the  
 446 construction of  $B(k)$  upon admission of some job  $k$  by job  $j$ , we also have  $B(k) \subseteq$   
 447  $S(j)$ .  $\square$

448 *Proof of Theorem 3.1.* Let  $j$  be a job admitted by the blocking algorithm to ma-  
 449 chine  $i$  with  $a_j \leq d_j - (1 + \delta)p_{ij}$ . Showing that job  $j$  completes before time  $d'_j :=$   
 450  $a_j + (1 + \delta)p_{ij}$  proves the theorem. Due to scheduling in SPT order, each job  $j$  has  
 451 highest priority in its own scheduling interval if the time point does not belong to the  
 452 scheduling interval of a descendant of  $j$ . Thus, it suffices to show that at most  $\delta p_{ij}$   
 453 units of time in  $[a_j, d'_j)$  belong to scheduling intervals  $S(k)$  of descendants of  $j$ . By  
 454 Lemma 3.2, the scheduling interval of any descendant  $k'$  of a child  $k$  of  $j$  is contained  
 455 in  $S(k)$ . Hence, it is sufficient to only consider  $K$ , the set of children of  $j$ .

456 In order to bound the contribution of each child  $k \in K$ , we impose a *class struc-*  
 457 *ture* on the jobs in  $K$  depending on their size relative to job  $j$ . More precisely, we  
 458 define  $(\mathcal{C}_c(j))_{c \in \mathbb{N}_0}$ , where  $\mathcal{C}_c(j)$  contains all jobs  $k \in K$  that satisfy  $\frac{\gamma}{2c+1}p_{ij} \leq p_{ik} <$   
 459  $\frac{\gamma}{2c}p_{ij}$ . As  $k \in K$  implies  $p_{ik} < \gamma p_{ij}$ , each child of  $j$  belongs to exactly one class  
 460 and  $(\mathcal{C}_c(j))_{c \in \mathbb{N}_0}$  in fact partitions  $K$ .

461 Consider two jobs  $k, k' \in K$  where, upon admission,  $k$  interrupts the blocking  
 462 period of  $k'$ . By definition, we have  $p_{ik} < \frac{1}{2}p_{ik'}$ . Hence, the chosen class structure  
 463 ensures that  $k$  belongs to a strictly *higher* class than  $k'$ , i.e., there are  $c, c' \in \mathbb{N}$   
 464 with  $c > c'$  such that  $k \in \mathcal{C}_c(j)$  and  $k' \in \mathcal{C}_{c'}(j)$ . In particular, the admission of a  
 465 job  $k \in \mathcal{C}_c(j)$  implies either that  $k$  is the first job of class  $\mathcal{C}_c(j)$  that  $j$  admits or that  
 466 the blocking period of the previous job in class  $\mathcal{C}_c(j)$  has completed. Based on this  
 467 distinction, we are able to bound the loss of scheduling time for  $j$  in  $S(j)$  due to  $S(k)$   
 468 of a child  $k$ . Specifically, we partition  $K$  into two sets. The first set  $K_1$  contains all  
 469 children of  $j$  that were admitted as the first jobs in their class  $\mathcal{C}_c(j)$ . The set  $K_2$   
 470 contains the remaining jobs.

471 We start with  $K_2$ . Consider a job  $k \in \mathcal{C}_c(j)$  admitted by  $j$ . By Lemma 3.2, we  
 472 know that  $|S(k)| = (1 + \mu\delta)p_{ik}$ , where  $1 \leq \mu \leq 2$ . Let  $k' \in \mathcal{C}_{c'}(j)$  be the previous job  
 473 admitted by  $j$  in class  $\mathcal{C}_{c'}(j)$ . Then,  $B(k') \subseteq [e_{k'}, a_k)$ . Since scheduling and blocking  
 474 periods of children of  $j$  are disjoint,  $j$  has highest scheduling priority in  $B(k')$ . Hence,  
 475 during  $B(k') \cup S(k)$  job  $j$  is processed for at least  $|B(k')|$  units of time. In other

476 words,  $j$  is processed for at least a  $\frac{|B(k')|}{|B(k') \cup S(k)|}$ -fraction of  $B(k') \cup S(k)$ . We rewrite  
 477 this ratio as

$$478 \quad \frac{|B(k')|}{|B(k') \cup S(k)|} = \frac{\beta p_{ik'}}{\beta p_{ik'} + (1 + \mu\delta)p_{ik}} = \frac{\nu\beta}{\nu\beta + (1 + \mu\delta)},$$

479 where  $\nu := \frac{p_{ik'}}{p_{ik}} \in (\frac{1}{2}, 2]$ . By differentiating with respect to  $\nu$  and  $\mu$ , we observe  
 480 that the last term is increasing in  $\nu$  and decreasing in  $\mu$ . Thus, we lower bound this  
 481 expression by

$$482 \quad \frac{|B(k')|}{|B(k') \cup S(k)|} \geq \frac{\beta/2}{\beta/2 + (1 + 2\delta)}.$$

483 Therefore,  $j$  is processed for at least a  $\frac{\beta/2}{\beta/2 + (1 + 2\delta)}$ -fraction in  $\bigcup_{k \in K} B(k) \cup \bigcup_{k \in K_2} S(k)$ .

484 We now consider the set  $K_1$ . The total processing volume of these jobs is bounded  
 485 from above by  $\sum_{c=0}^{\infty} \frac{\gamma}{2^c} p_{ij} = 2\gamma p_{ij}$ . By [Lemma 3.2](#),  $|S(k)| \leq (1 + 2\delta)p_{ik}$ . Combining  
 486 these two observations, we obtain  $|\bigcup_{k \in K_1} S(k)| \leq 2(1 + 2\delta)\gamma p_{ij}$ . Combining the latter  
 487 with the bound for  $K_2$ , we conclude that  $j$  is scheduled for at least

$$488 \quad \left| [a_j, d'_j] \setminus \bigcup_{k \in K} S(k) \right| \geq \frac{\beta/2}{\beta/2 + (1 + 2\delta)} ((1 + \delta) - 2(1 + 2\delta)\gamma) p_{ij} \geq p_{ij}$$

489 units of time, where the last inequality follows from Equation (3.1). Therefore,  $j$   
 490 completes before  $d'_j = a_j + (1 + \delta)p_{ij} \leq d_j$ , which concludes the proof.  $\square$

491 **4. Competitiveness: admitting sufficiently many jobs.** This section shows  
 492 that the blocking algorithm admits sufficiently many jobs to be  $\mathcal{O}(\frac{1}{\varepsilon - \delta})$ -competitive.  
 493 As mentioned before, this proof is based on the observation that, at time  $\tau$ , the  
 494 blocking algorithm admits any job available for machine  $i$  if its processing time is  
 495 less than  $\gamma p_{ij}$ , where  $j_i$  is the job processed by machine  $i$  at time  $\tau$ , and this time  
 496 point is not blocked by another job  $k_i$  previously admitted by  $j_i$  to machine  $i$ . We  
 497 start by formalizing this observation for a class of non-migratory online algorithms  
 498 before proving that this enables us to bound the number of jobs any feasible schedule  
 499 successfully schedules during a particular period. Then, we use it to show that the  
 500 blocking algorithm is indeed  $\mathcal{O}(\frac{1}{\varepsilon - \delta})$ -competitive.

501 **4.1. A class of online algorithms.** In this section, we investigate a class of  
 502 non-migratory online algorithms. Recall that a job  $j$  is called available for machine  $i$   
 503 at time  $\tau$  if it is released before or at time  $\tau$ ,  $d_j - \tau \geq (1 + \delta)p_{ij}$ , and is not yet  
 504 admitted.

505 We consider a non-migratory online algorithm  $\mathcal{A}$  with the following properties.

506 (P1)  $\mathcal{A}$  only admits available jobs.

507 (P2) Retrospectively, for each time  $\tau$  and each machine  $i$ , there is a threshold  $u_{i\tau} \in$   
 508  $[0, \infty]$  such that any job  $j$  that was available for machine  $i$  and not admit-  
 509 ted to machine  $i$  by  $\mathcal{A}$  at time  $\tau$  satisfies  $p_{ij} \geq u_{i\tau}$ . The function  $u^{(i)} :$   
 510  $\mathbb{R} \rightarrow [0, \infty], \tau \mapsto u_{i\tau}$  is piece-wise constant and right-continuous for every  
 511 machine  $i \in \{1, \dots, m\}$ . Further, there are only countably many points of  
 512 discontinuity. (This last property is used to simplify the exposition.)

513 **Key lemma on the size of non-admitted jobs.** For the proof of the main  
 514 result in this section, we rely on the following strong, structural lemma about the  
 515 volume processed by a feasible non-migratory schedule in some time interval and the

516 size of jobs admitted by a non-migratory online algorithm satisfying (P1) and (P2) in  
 517 the same time interval.

518 Let  $\sigma$  be a feasible non-migratory schedule. Without loss of generality, we assume  
 519 that  $\sigma$  completes all jobs that it started on time. Let  $X^\sigma$  be the set of jobs completed  
 520 by  $\sigma$  and not admitted by  $\mathcal{A}$ . For  $1 \leq i \leq m$ , let  $X_i^\sigma$  be the set of jobs in  $X^\sigma$  processed  
 521 by machine  $i$ . Let  $C_x$  be the completion time of job  $x \in X^\sigma$  in  $\sigma$ .

522 LEMMA 4.1. *Let  $0 \leq \vartheta_1 \leq \vartheta_2$  and fix  $x \in X_i^\sigma$  as well as  $Y \subset X_i^\sigma \setminus \{x\}$ . If*

523 *(R)  $r_x \geq \vartheta_1$  as well as  $r_y \geq \vartheta_1$  for all  $y \in Y$ ,*

524 *(C)  $C_x \geq C_y$  for all  $y \in Y$ , and*

525 *(P)  $\sum_{y \in Y} p_{iy} \geq \frac{\varepsilon}{\varepsilon - \delta}(\vartheta_2 - \vartheta_1)$*

526 *hold, then  $p_{ix} \geq u_{i\vartheta_2}$ , where  $u_{i\vartheta_2}$  is the threshold imposed by  $\mathcal{A}$  at time  $\vartheta_2$ . In*  
 527 *particular, if  $u_{i,\vartheta_2} = \infty$ , then no such job  $x$  exists.*

528 *Proof.* We show the lemma by contradiction. More precisely, we show that,  
 529 if  $p_{ix} < u_{i\vartheta_2}$ , the schedule  $\sigma$  cannot complete  $x$  on time and, hence, is not feasi-  
 530 ble.

531 Remember that  $x \in X_i^\sigma$  implies that  $\mathcal{A}$  did not admit job  $x$  at any point  $\vartheta$ .  
 532 At time  $\vartheta_2$ , there are two possible reasons why  $x$  was not admitted:  $p_{ix} \geq u_{i\vartheta_2}$  or  
 533  $d_x - \vartheta_2 < (1 + \delta)p_{ix}$ . In case of the former, the statement of the lemma holds. Toward  
 534 a contradiction, suppose  $p_{ix} < u_{i\vartheta_2}$  and, thus,  $d_x - \vartheta_2 < (1 + \delta)p_{ix}$  has to hold.  
 535 As job  $x$  arrives with a slack of at least  $\varepsilon p_{ix}$  at its release date  $r_x$  and  $r_x \geq \vartheta_1$  by  
 536 assumption, we have

$$537 \quad (4.1) \quad \vartheta_2 - \vartheta_1 \geq \vartheta_2 - d_x + d_x - r_x > -(1 + \delta)p_{ix} + (1 + \varepsilon)p_{ix} = (\varepsilon - \delta)p_{ix}.$$

538 Since all jobs in  $Y$  complete earlier than  $x$  by Assumption (C) and are only  
 539 released after  $\vartheta_1$  by (R), the volume processed by  $\sigma$  in  $[\vartheta_1, C_x)$  on machine  $i$  is at  
 540 least  $\frac{\varepsilon}{\varepsilon - \delta}(\vartheta_2 - \vartheta_1) + p_{ix}$  by (P). Moreover,  $\sigma$  can process at most a volume of  $(\vartheta_2 - \vartheta_1)$   
 541 on machine  $i$  in  $[\vartheta_1, \vartheta_2)$ . These two bounds imply that  $\sigma$  has to process job parts  
 542 with a processing volume of at least

$$543 \quad \frac{\varepsilon}{\varepsilon - \delta}(\vartheta_2 - \vartheta_1) + p_{ix} - (\vartheta_2 - \vartheta_1) > \frac{\delta}{\varepsilon - \delta}(\varepsilon - \delta)p_{ix} + p_{ix} = (1 + \delta)p_{ix}$$

544 in  $[\vartheta_2, C_x)$ , where the inequality follows using Inequality (4.1). Thus,  $C_x \geq \vartheta_2 + (1 +$   
 545  $\delta)p_{ix} > d_x$ , which contradicts the feasibility of  $\sigma$ .

546 Observe that, by (P1) and (P2), the online algorithm  $\mathcal{A}$  admits a job available  
 547 for machine  $i$  if it satisfies  $p_{ij} < u_{i\tau}$ . In particular, if  $u_{i\tau} = \infty$  for some time point  $\tau$ ,  
 548 then  $\mathcal{A}$  admits any job available for machine  $i$ . Hence, for  $0 \leq \vartheta_1 \leq \vartheta_2$  with  $u_{i\vartheta_2} = \infty$ ,  
 549 there does not exist a job  $x \in X_i^\sigma$  and a set  $Y \subset X_i^\sigma \setminus \{x\}$  satisfying (R), (C), and  
 550 (P) for machine  $i$ .  $\square$

551 **Bounding the number of non-admitted jobs.** In this section, we use the  
 552 Properties (P1) and (P2) to bound the throughput of a non-migratory optimal (offline)  
 553 algorithm. To this end, we fix an instance as well as an optimal schedule with job set  
 554 OPT. Let  $\mathcal{A}$  be a non-migratory online algorithm satisfying (P1) and (P2).

555 Let  $X$  be the set of jobs in OPT that the algorithm  $\mathcal{A}$  did not admit. We assume  
 556 without loss of generality that all jobs in OPT complete on time. Since OPT as well  
 557 as  $\mathcal{A}$  are non-migratory, we compare the throughput machine-wise. To this end, we  
 558 fix one machine  $i$ . Let  $X_i \subset X$  be the set of jobs scheduled on machine  $i$  by OPT.

559 Assumption (P2) guarantees that the threshold  $u_{i,\tau}$  is piece-wise constant and  
 560 right-continuous, i.e.,  $u^{(i)}$  is constant on intervals of the form  $[\tau_t, \tau_{t+1})$ . Let  $\mathcal{I}$  represent

561 the set of maximal intervals  $I_t = [\tau_t, \tau_{t+1})$  where  $u^{(i)}$  is constant. That is,  $u_{i,\tau} = u_t$   
 562 holds for all  $\tau \in I_t$  and  $u_{i,\tau_{t+1}} \neq u_t$ , where  $u_t := u_{i,\tau_t}$ . The main result of this section  
 563 is the following theorem.

564 **THEOREM 4.2.** *Let  $X_i$  be the set of jobs that are scheduled on machine  $i$  in the*  
 565 *optimal schedule. Let  $\mathcal{I} = \{I_1, \dots, I_T\}$  be the set of maximal intervals on machine  $i$*   
 566 *of  $\mathcal{A}$  such that the machine-dependent threshold is constant for each interval and has*  
 567 *the value  $u_t$  in interval  $I_t = [\tau_t, \tau_{t+1})$ . Then,*

$$568 \quad |X_i| \leq \sum_{t=1}^T \frac{\varepsilon}{\varepsilon - \delta} \frac{\tau_{t+1} - \tau_t}{u_t} + T,$$

569 where we set  $\frac{\tau_{t+1} - \tau_t}{u_t} = 0$  if  $u_t = \infty$  and  $\frac{\tau_{t+1} - \tau_t}{u_t} = \infty$  if  $\{\tau_t, \tau_{t+1}\} \cap \{-\infty, \infty\} \neq \emptyset$   
 570 and  $u_t < \infty$ .

571 We observe that  $T = \infty$  trivially proves the statement as  $X_i$  contains at most  
 572 finitely many jobs. The same is true if  $\frac{\tau_{t+1} - \tau_t}{u_t} = \infty$  for some  $t \in [T]$ . Hence, for the  
 573 remainder of this section we assume without loss of generality that  $\mathcal{I}$  only contains  
 574 finitely many intervals and that  $\frac{\tau_{t+1} - \tau_t}{u_t} < \infty$  holds for every  $t \in [T]$ .

575 To prove this theorem, we develop a charging scheme that assigns jobs  $x \in X_i$   
 576 to intervals in  $\mathcal{I}$ . The idea behind our charging scheme is that OPT does not contain  
 577 arbitrarily many jobs that are available in  $I_t$  since  $u_t$  provides a natural lower bound on  
 578 their processing times. In particular, the processing time of any job that is *released*  
 579 during interval  $I_t$  and not admitted by the algorithm exceeds the lower bound  $u_t$ .  
 580 Thus, the charging scheme relies on the release date  $r_x$  and the size  $p_{ix}$  of a job  $x \in X_i$   
 581 as well as on the precise structure of the intervals created by  $\mathcal{A}$ .

582 The charging scheme we develop is based on a careful modification of the following  
 583 partition  $(F_t)_{t=1}^T$  of the set  $X_i$ . Fix an interval  $I_t \in \mathcal{I}$  and define the set  $F_t \subseteq X_i$  as  
 584 the set that contains all jobs  $x \in X_i$  released during  $I_t$ , i.e.,  $F_t = \{x \in X_i : r_x \in I_t\}$ .  
 585 Since, upon release, each job  $x \in X_i$  is available and not admitted by  $\mathcal{A}$ , the next fact  
 586 directly follows from Properties (P1) and (P2).

587 **FACT 4.3.** *For all jobs  $x \in F_t$  it holds  $p_{ix} \geq u_t$ . In particular, if  $u_t = \infty$ ,*  
 588 *then  $F_t = \emptyset$ .*

589 In fact, the charging scheme maintains this property and only assigns jobs in  $X_i$   
 590 to intervals  $I_t$  if  $p_{ix} \geq u_t$ . In particular, no job will be assigned to an interval  
 591 with  $u_t = \infty$ .

592 We now formalize how many jobs in  $X_i$  are assigned to a specific interval  $I_t$ . Let

$$593 \quad \varphi_t := \left\lfloor \frac{\varepsilon}{\varepsilon - \delta} \frac{\tau_{t+1} - \tau_t}{u_t} \right\rfloor + 1$$

594 if  $u_t < \infty$ , and  $\varphi_t = 0$  if  $u_t = \infty$ . We refer to  $\varphi_t$  as the *target number* of  $I_t$ . As  
 595 discussed before, we assume  $\frac{\tau_{t+1} - \tau_t}{u_t} < \infty$ , and, thus, the target number is well-defined.  
 596 If each of the sets  $F_t$  satisfies  $|F_t| \leq \varphi_t$ , then **Theorem 4.2** immediately follows. In  
 597 general,  $|F_t| \leq \varphi_t$  does not have to be true since jobs in OPT may be preempted and  
 598 processed during several intervals  $I_t$ . Therefore, for proving **Theorem 4.2**, we show  
 599 that there always exists another partition  $(G_t)_{t=1}^T$  of  $X_i$  such that  $|G_t| \leq \varphi_t$  holds.

600 The high-level idea of this proof is the following: Consider an interval  $I_t =$   
 601  $[\tau_t, \tau_{t+1})$ . If  $F_t$  does not contain too many jobs, i.e.,  $|F_t| \leq \varphi_t$ , we would like to  
 602 set  $G_t = F_t$ . Otherwise, we find a later interval  $I_{t'}$  with  $|F_{t'}| < \varphi_{t'}$  such that we can  
 603 assign the excess jobs in  $F_t$  to  $I_{t'}$ .

604 *Proof of Theorem 4.2.* As observed before, it suffices to show the existence of a  
 605 partition  $\mathcal{G} = (G_t)_{t=1}^T$  of  $X_i$  such that  $|G_t| \leq \varphi_t$  in order to prove the theorem.

606 In order to repeatedly apply Lemma 4.1, we only assign excess jobs  $x \in F_t$  to  $G_{t'}$   
 607 for  $t < t'$  if their processing time is at least the threshold of  $I_{t'}$ , i.e.,  $p_{ix} \geq u_{t'}$ . By  
 608 our choice of parameters, a set  $G_{t'}$  with  $\varphi_{t'}$  many jobs of size at least  $u_{t'}$  “covers” the  
 609 interval  $I_{t'} = [\tau_{t'}, \tau_{t'+1})$  as often as required by (P) in Lemma 4.1, i.e.,

$$610 \quad (4.2) \quad \sum_{x \in G_{t'}} p_{ix} \geq \varphi_{t'} \cdot u_{t'} = \left( \left\lfloor \frac{\varepsilon}{\varepsilon - \delta} \frac{\tau_{t'+1} - \tau_{t'}}{u_{t'}} \right\rfloor + 1 \right) \cdot u_{t'} \geq \frac{\varepsilon}{\varepsilon - \delta} (\tau_{t'+1} - \tau_{t'}).$$

611 The proof consists of two parts: the first one is to inductively (on  $t$ ) construct the  
 612 partition  $\mathcal{G} = (G_t)_{t=1}^T$  of  $X_i$ , where  $|G_t| \leq \varphi_t$  holds for  $t \in [T - 1]$ . The second one  
 613 is the proof that a job  $x \in G_t$  satisfies  $p_{ix} \geq u_t$  which will imply  $|G_T| \leq \varphi_T$ . During  
 614 the construction of  $\mathcal{G}$  we define temporary sets  $A_t \subset X_i$  for intervals  $I_t$ . The set  $G_t$   
 615 is chosen as a subset of  $F_t \cup A_t$  of appropriate size. In order to apply Lemma 4.1 to  
 616 each job in  $A_t$  individually, alongside  $A_t$ , we construct a set  $Y_{x,t}$  and a time  $\tau_{x,t} \leq r_x$   
 617 for each job  $x \in X_i$  that is added to  $A_t$ . Let  $C_y^*$  be the completion time of some  
 618 job  $y \in X_i$  in the optimal schedule OPT. The second part of the proof is to show  
 619 that  $x$ ,  $\tau_{x,t}$ , and  $Y_{x,t}$  satisfy

- 620 (R)  $r_y \geq \tau_{x,t}$  for all  $y \in Y_{x,t}$ ,
- 621 (C)  $C_x^* \geq C_y^*$  for all  $y \in Y_{x,t}$ , and
- 622 (P)  $\sum_{y \in Y_{x,t}} p_{iy} \geq \frac{\varepsilon}{\varepsilon - \delta} (\tau_t - \tau_{x,t})$ .

623 This implies that  $x$ ,  $Y = Y_{x,t}$ ,  $\vartheta_1 = \tau_{x,t}$ , and  $\vartheta_2 = \tau_t$  satisfy the conditions of  
 624 Lemma 4.1, and thus the processing time of  $x$  is at least the threshold at time  $\tau_t$ ,  
 625 i.e.,  $p_{ix} \geq u_{i\tau_t} = u_t$ .

626 *Constructing  $G = (G_t)_{t=1}^T$ .* We inductively construct the sets  $G_t$  in the order  
 627 of their indices. We start by setting  $A_t = \emptyset$  for all intervals  $I_t$  with  $t \in T$ . We  
 628 define  $Y_{x,t} = \emptyset$  for each job  $x \in X_i$  and each interval  $I_t$ . The preliminary value of  
 629 the time  $\tau_{x,t}$  is the minimum of the starting point  $\tau_t$  of the interval  $I_t$  and the release  
 630 date  $r_x$  of  $x$ , i.e.,  $\tau_{x,t} := \min\{\tau_t, r_x\}$ . We refer to the step in the construction where  $G_t$   
 631 was defined by *step t*.

632 Starting with  $t = 1$ , let  $I_t$  be the next interval to consider during the construction  
 633 with  $t < T$ . Depending on the cardinality of  $F_t \cup A_t$ , we distinguish two cases. If  
 634  $|F_t \cup A_t| \leq \varphi_t$ , then we set  $G_t = F_t \cup A_t$ .

635 If  $|F_t \cup A_t| > \varphi_t$ , then we order the jobs in  $F_t \cup A_t$  in increasing order of com-  
 636 pletion times in the optimal schedule. The first  $\varphi_t$  jobs are assigned to  $G_t$  while the  
 637 remaining  $|F_t \cup A_t| - \varphi_t$  jobs are added to  $A_{t+1}$ . In this case, we might have to  
 638 redefine the times  $\tau_{x,t+1}$  and the sets  $Y_{x,t+1}$  for the jobs  $x$  that were newly added  
 639 to  $A_{t+1}$ . Fix such a job  $x$ . If there is no job  $z$  in the just defined set  $G_t$  that has a  
 640 smaller release date than  $\tau_{x,t}$ , we set  $\tau_{x,t+1} = \tau_{x,t}$  and  $Y_{x,t+1} = Y_{x,t} \cup G_t$ . Otherwise  
 641 let  $z \in G_t$  be a job with  $r_z < \tau_{x,t}$  that has the smallest time  $\tau_{z,t}$ . We set  $\tau_{x,t+1} = \tau_{z,t}$   
 642 and  $Y_{x,t+1} = Y_{z,t} \cup G_t$ .

643 Finally, we set  $G_T = F_T \cup A_T$ . We observe that  $u_T < \infty$  implies  $\varphi_T = \infty$   
 644 because  $\tau_{T+1} = \infty$ . Since this contradicts the assumption  $\varphi_t < \infty$  for all  $t \in [T]$ ,  
 645 this implies  $u_T = \infty$ . We will show that  $p_x \geq u_T$  for all  $x \in G_T$ . Hence,  $G_T = \emptyset$ .  
 646 Therefore  $|G_T| = \varphi_T = 0$ .

647 *Bounding the size of jobs in  $G_t$ .* We consider the intervals again in increasing  
 648 order of their indices and show by induction that any job  $x$  in  $G_t$  satisfies  $p_{ix} \geq u_t$   
 649 which implies  $G_t = \emptyset$  if  $u_t = \infty$ . Clearly, if  $x \in F_t \cap G_t$ , Fact 4.3 guarantees  $p_{ix} \geq u_t$ .  
 650 Hence, in order to show the lower bound on the processing time of  $x \in G_t$ , it is



sufficient to consider jobs in  $G_t \setminus F_t \subset A_t$ . To this end, we show that for such jobs (R), (C), and (P) are satisfied. Thus, Lemma 4.1 guarantees that  $p_{ix} \geq u_{i\tau_t} = u_t$  by definition. Hence,  $A_t = \emptyset$  if  $u_t = \infty$  by Lemma 4.1.

By construction,  $A_1 = \emptyset$ . Hence, (R), (C), and (P) are satisfied for each job  $x \in A_1$ .

Suppose that the Conditions (R), (C), and (P) are satisfied for all  $x \in A_s$  for all  $1 \leq s < t$ . Hence, for  $s < t$ , the set  $G_s$  only contains jobs  $x$  with  $p_{ix} \geq u_s$ . Fix  $x \in A_t$ . We want to show that  $p_{ix} \geq u_t$ . By the induction hypothesis and by Fact 4.3,  $p_{iy} \geq u_{t-1}$  holds for all  $y \in G_{t-1}$ . Since  $x$  did not fit in  $G_{t-1}$  anymore,  $|G_{t-1}| = \varphi_{t-1}$ .

We distinguish two cases based on  $G_{t-1}$ . If there is no job  $z \in G_{t-1}$  with  $r_z < \tau_{x,t-1}$ , then  $\tau_{x,t} = \tau_{x,t-1}$ , and (R) and (C) are satisfied by construction and by the induction hypothesis. For (P), consider

$$\begin{aligned} \sum_{y \in Y_{x,t}} p_{iy} &= \sum_{y \in Y_{x,t-1}} p_{iy} + \sum_{y \in G_{t-1}} p_{iy} \\ &\geq \frac{\varepsilon}{\varepsilon - \delta} (\tau_{t-1} - \tau_{x,t-1}) + u_{t-1} \cdot \varphi_{t-1} \\ &\geq \frac{\varepsilon}{\varepsilon - \delta} (\tau_{t-1} - \tau_{x,t-1}) + \frac{\varepsilon}{\varepsilon - \delta} (\tau_t - \tau_{t-1}) \\ &= \frac{\varepsilon}{\varepsilon - \delta} (\tau_t - \tau_{x,t}), \end{aligned}$$

where the first inequality holds due to the induction hypothesis. By Lemma 4.1,  $p_{ix} \geq u_{\tau_t} = u_t$ .

If there is a job  $z \in G_{t-1}$  with  $r_z < \tau_{x,t-1} \leq \tau_{t-1}$ , then  $z \in A_{t-1}$ . In step  $t-1$ , we chose  $z$  with minimal  $\tau_{z,t-1}$ . Thus,  $r_y \geq \tau_{y,t-1} \geq \tau_{z,t-1}$  for all  $y \in G_{t-1}$  and  $r_x \geq \tau_{x,t-1} > r_z \geq \tau_{z,t-1}$  which is Condition (R) for the jobs in  $G_{t-1}$ . Moreover, by the induction hypothesis,  $r_y \geq \tau_{z,t-1}$  holds for all  $y \in Y_{z,t-1}$ . Thus,  $\tau_{x,t}$  and  $Y_{x,t}$  satisfy (R). For (C), consider that  $C_x^* \geq C_y^*$  for all  $y \in G_{t-1}$  by construction and, thus,  $C_x^* \geq C_z^* \geq C_y^*$  also holds for all  $y \in Y_{z,t-1}$  due to the induction hypothesis. For (P), observe that

$$\begin{aligned} \sum_{y \in Y_{x,t}} p_{iy} &= \sum_{y \in Y_{z,t-1}} p_{iy} + \sum_{y \in G_{t-1}} p_{iy} \\ &\geq \frac{\varepsilon}{\varepsilon - \delta} (\tau_{t-1} - \tau_{z,t-1}) + u_{t-1} \cdot \varphi_{t-1} \\ &\geq \frac{\varepsilon}{\varepsilon - \delta} (\tau_{t-1} - \tau_{z,t-1}) + \frac{\varepsilon}{\varepsilon - \delta} (\tau_t - \tau_{t-1}) \\ &\geq \frac{\varepsilon}{\varepsilon - \delta} (\tau_t - \tau_{x,t}). \end{aligned}$$

Here, the first inequality follows from the induction hypothesis and the second from the definition of  $u_{t-1}$  and  $\varphi_{t-1}$ . Hence, Lemma 4.1 implies  $p_{ix} \geq u_{\tau_t} = u_t$ .

We note that  $p_{ix} \geq u_t$  for all  $x \in G_t$  and for all  $t \in [T]$ .

*Bounding  $|X_i|$ .* By construction, we know that  $\bigcup_{t=1}^T G_t = X_i$ . We start with considering intervals  $I_t$  with  $u_t = \infty$ . Then,  $I_t$  has an unbounded threshold, i.e.,  $u_{i\tau} = \infty$  for all  $\tau \in I_t$ , and  $F_t = \emptyset$  by Fact 4.3. In the previous part we have seen that the conditions for Lemma 4.1 are satisfied. Hence,  $G_t = \emptyset$  if  $u_t = \infty$ . For  $t$  with  $u_t < \infty$ , we have  $|G_t| \leq \varphi_t = \lfloor \frac{\varepsilon}{\varepsilon - \delta} \frac{\tau_{t+1} - \tau_t}{u_t} \rfloor + 1$ . As explained before, this bounds the number of jobs in  $X_i$ .  $\square$

692 **4.2. The blocking algorithm admits sufficiently many jobs.** Having the  
 693 powerful tool that we developed in the previous section at hand, it remains to show  
 694 that the blocking algorithm admits sufficiently many jobs to achieve the competitive  
 695 ratio of  $\mathcal{O}\left(\frac{1}{\varepsilon-\delta'}\right)$  where  $\delta' = \frac{\varepsilon}{2}$  for commitment upon admission and  $\delta' = \max\left\{\frac{\varepsilon}{2}, \delta\right\}$   
 696 for  $\delta$ -commitment. To this end, we show that the blocking algorithm belongs to the  
 697 class of online algorithms considered in [Subsection 4.1](#). Then, [Theorem 4.2](#) provides  
 698 a bound on the throughput of an optimal non-migratory schedule.

699 We begin by showing that the blocking algorithm satisfies Properties (P1) to (P2).  
 700 The first property is clearly satisfied by the definition of the blocking algorithm. For  
 701 the second and the third property, we observe that a new job  $j^*$  is only admitted  
 702 to a machine  $i$  during the scheduling interval of another job  $j$  admitted to the same  
 703 machine if  $p_{ij^*} < \gamma p_{ij}$ . Further, the time point of admission must not be blocked by a  
 704 similar- or smaller-size job  $k$  previously admitted during the scheduling interval of  $j$ .  
 705 This leads to the bound  $p_{ij^*} < \frac{1}{2}p_{ik}$  for any job  $k$  whose blocking period contains  
 706 the current time point. Combining these observations leads to a machine-dependent  
 707 threshold  $u_{i,\tau} \in [0, \infty]$  satisfying (P2).

708 More precisely, fix a machine  $i$  and a time point  $\tau$ . Using  $j \rightarrow i$  to denote that  $j$   
 709 was admitted to machine  $i$ , we define  $u_{i,\tau} := \min_{j: j \rightarrow i, \tau \in S(j)} \gamma p_{ij}$  if there is no job  $k$   
 710 admitted to machine  $i$  with  $\tau \in B(k)$ , with  $\min \emptyset = \infty$ . Otherwise, we set  $u_{i,\tau} := \frac{1}{2}p_{ik}$ .  
 711 We note that the function  $u^{(i)}$  is piece-wise constant and right-continuous due to our  
 712 choice of right-open intervals for defining scheduling intervals and blocking periods.  
 713 Moreover, the points of discontinuity of  $u^{(i)}$  correspond to the admission of a new job,  
 714 the end of a scheduling interval, and the start as well as the end of a blocking period  
 715 of jobs admitted to machine  $i$ . Since we only consider instances with a finite number  
 716 of jobs, there are at most finitely many points of discontinuity of  $u^{(i)}$ . Hence, we can  
 717 indeed apply [Theorem 4.2](#).

718 Then, the following theorem is the main result of this section.

719 **THEOREM 4.4.** *An optimal non-migratory (offline) algorithm can complete at*  
 720 *most a factor  $\alpha + 5$  more jobs on time than admitted by the blocking algorithm, where*  
 721  $\alpha := \frac{\varepsilon}{\varepsilon-\delta} \left(2\beta + \frac{1+2\delta}{\gamma}\right)$ .

722 *Proof.* We fix an instance and an optimal solution OPT. We use  $X$  to denote  
 723 the set of jobs in OPT that the blocking algorithm did not admit. Without loss of  
 724 generality, we can assume that all jobs in OPT complete on time. If  $J$  is the set of jobs  
 725 admitted by the blocking algorithm, then  $X \cup J$  is a superset of the jobs successfully  
 726 finished in the optimal solution. Hence, showing  $|X| \leq (\alpha + 4)|J|$  suffices to prove  
 727 [Theorem 4.4](#).

728 For each machine  $i$ , we compare the throughput of the optimal solution to the  
 729 throughput on machine  $i$  of the blocking algorithm. More precisely, let  $X_i \subseteq X$  be  
 730 the jobs in OPT scheduled on machine  $i$  and let  $J_i \subseteq J$  be the jobs scheduled by the  
 731 blocking algorithm on machine  $i$ . With [Theorem 4.2](#), we show  $|X_i| \leq (\alpha + 4)|J_i|$  to  
 732 bound the cardinality of  $X$  in terms of  $|J|$ .

733 To this end, we retrospectively consider the interval structure created by the  
 734 algorithm on machine  $i$ . Let  $\mathcal{I}$  be the set of maximal intervals  $I_t = [\tau_t, \tau_{t+1})$  such  
 735 that  $u_{i,\tau} = u_{i,\tau_t}$  for all  $\tau \in I_t$ . We define  $u_t = u_{i,\tau_t}$  for each interval  $I_t$ . As discussed  
 736 above, the time points  $\tau_t$  for  $t \in [T]$  correspond to the admission, the end of a  
 737 scheduling interval, and the start as well as the end of a blocking period of jobs  
 738 admitted to machine 1. As the admission of a job adds at most three time points, we  
 739 have that  $|\mathcal{I}| \leq 3|J_i| + 1$ .

740 As the blocking algorithm satisfies Properties (P1) to (P2), we can apply [Theo-](#)

741 rem 4.2 to obtain

$$742 \quad |X_i| \leq \sum_{t=1}^T \frac{\varepsilon}{\varepsilon - \delta} \frac{\tau_{t+1} - \tau_t}{u_t} + |\mathcal{I}| \leq \sum_{t=1}^T \frac{\varepsilon}{\varepsilon - \delta} \frac{\tau_{t+1} - \tau_t}{u_t} + (3|J_i| + 1).$$

743 It remains to bound the first part in terms of  $|J_i|$ . If  $u_t < \infty$ , let  $j_t \in J_i$  be the *smallest*  
 744 job  $j$  with  $\tau_t \in S(j) \cup B(j)$ . Then, at most  $\frac{\varepsilon}{\varepsilon - \delta} \frac{\tau_{t+1} - \tau_t}{u_t}$  (potentially fractional) jobs  
 745 will be charged to job  $j_t$  because of interval  $I_t$ . By definition of  $u_t$ , we have  $u_t = \gamma p_{ij_t}$   
 746 if  $I_t \subseteq S(j_t)$ , and if  $I_t \subseteq B(j_t)$ , we have  $u_t = \frac{1}{2} p_{ij_t}$ . The total length of intervals  $I_t$  for  
 747 which  $j = j_t$  holds sums up to at most  $(1 + 2\delta)p_{ij}$  for  $I_t \subseteq S(j)$  and to at most  $2\beta p_{ij}$   
 748 for  $I_t \subseteq B(j)$ . Hence, in total, the charging scheme assigns at most  $\frac{\varepsilon}{\varepsilon - \delta} (2\beta + \frac{1+2\delta}{\gamma}) = \alpha$   
 749 jobs in  $X_i$  to job  $j \in J_i$ . Therefore,

$$750 \quad |X_i| \leq (\alpha + 3)|J_i| + 1.$$

751 If  $J_i = \emptyset$ , the blocking algorithm admitted all jobs scheduled on machine  $i$  by OPT,  
 752 and  $|X_i| = 0 = |J_i|$  follows. Otherwise,  $|X_i| \leq (\alpha + 4)|J_i|$ , and we obtain

$$753 \quad |\text{OPT}| \leq |X \cup J| = \sum_{i=1}^m |X_i| + |J| \leq \sum_{i=1}^m (\alpha + 4)|J_i| + |J| \leq (\alpha + 5)|J|,$$

754 which concludes the proof.  $\square$

### 755 4.3. Finalizing the proof of Theorem 1.1.

756 *Proof of Theorem 1.1.* In Theorem 3.1 we show that the blocking algorithm com-  
 757 pletes all admitted jobs  $J$  on time. This implies that the blocking algorithm is  
 758 feasible for the model commitment upon admission. As no job  $j \in J$  is admit-  
 759 ted later than  $d_j - (1 + \delta)p_{ij}$ , the blocking algorithm also solves scheduling with  $\delta$ -  
 760 commitment. In Theorem 4.4, we bound the throughput  $|\text{OPT}|$  of an optimal non-  
 761 migratory solution by  $\alpha + 5$  times  $|J|$ , the throughput of the blocking algorithm, where  
 762  $\alpha = \frac{\varepsilon}{\varepsilon - \delta} (2\beta + \frac{1+2\delta}{\gamma})$ . Our choice of parameters  $\beta = \frac{16}{\delta}$  and  $\gamma = \frac{\delta}{16}$  implies that the  
 763 blocking algorithm achieves a competitive ratio of  $c \in \mathcal{O}(\frac{\varepsilon}{(\varepsilon - \delta)\delta})$ . For commitment  
 764 upon arrival or for  $\delta$ -commitment in the case where  $\delta \leq \frac{\varepsilon}{2}$ , we run the algorithm  
 765 with  $\delta' = \frac{\varepsilon}{2}$ . Hence,  $c \in \mathcal{O}(\frac{1}{\varepsilon - \delta'}) = \mathcal{O}(\frac{1}{\varepsilon})$ . If  $\delta > \frac{\varepsilon}{2}$ , then we set  $\delta' = \delta$  in our  
 766 algorithm. Thus,  $\frac{\varepsilon}{\delta'} \in \mathcal{O}(1)$  and, again,  $c \in \mathcal{O}(\frac{1}{\varepsilon - \delta'})$ .  $\square$

767 **5. Scheduling without commitment.** This section considers online through-  
 768 put maximization without commitment requirements. We show how to exploit also  
 769 in this setting our key lemma on the size of non-admitted jobs for a big class of online  
 770 algorithms and the resulting upper bound on the throughput of an optimal (offline)  
 771 non-migratory algorithm from Subsection 4.1.

772 We consider the *region algorithm* that was designed by [10] for scheduling on a  
 773 single machine and we generalize it to parallel identical machines. We prove that  
 774 it has a competitive ratio of  $\mathcal{O}(\frac{1}{\varepsilon})$ , which is best possible on a single machine and  
 775 improves substantially upon the best previously known parallel-machine algorithm  
 776 (for weighted throughput) with a competitive ratio of  $\mathcal{O}(\frac{1}{\varepsilon^2})$  by Lucier et al. [31].  
 777 For a single machine, this matches the guarantee proven in [10]. However, our new  
 778 analysis is much more direct.

779 **5.1. The region algorithm.** Originally, the region algorithm was designed for  
 780 online scheduling with and without commitment on a single machine. We extend it to

781 unrelated machines by never migrating jobs between machines and per machine using  
 782 the same design principles that guide the admission decisions of the region algorithm,  
 783 as developed in [10]. Since we do not consider commitment in this section, we can  
 784 significantly simplify the exposition of the region algorithm when compared to [10].

785 As in the previous section, a job is only processed by the machine it initially was  
 786 started on. We say the job has been *admitted* to this machine. Moreover, a running  
 787 job can only be preempted by significantly smaller-size jobs, i.e., smaller by a factor  
 788 of at least  $\frac{\varepsilon}{4}$  with respect to the processing time, and a job  $j$  cannot start for the first  
 789 time on machine  $i$  when its remaining slack is too small, i.e., less than  $\frac{\varepsilon}{2}p_{ij}$ .

790 Formally, at any time  $\tau$ , the region algorithm maintains two sets of jobs: *admitted*  
 791 *jobs*, which have been started before or at time  $\tau$ , and *available jobs*. A job  $j$  is  
 792 available for machine  $i$  if it is released before or at time  $\tau$ , is not yet admitted, and  $\tau$   
 793 is not too close to its deadline, i.e.,  $r_j \leq \tau$  and  $d_j - \tau \geq (1 + \frac{\varepsilon}{2})p_{ij}$ . The intelligence  
 794 of the region algorithm lies in how it admits jobs. The actual scheduling decision  
 795 then is simple and independent of the admission of jobs: at any point in time and on  
 796 each machine, schedule the shortest job that has been admitted to this machine and  
 797 has not completed its processing time. In other words, we schedule admitted jobs on  
 798 each machine in SHORTEST PROCESSING TIME (SPT) order. The region algorithm  
 799 never explicitly considers deadlines except when deciding whether to admit jobs. In  
 800 particular, jobs can even be processed after their deadline.

801 At any time  $\tau$ , when there is a job  $j$  available for an *idle* machine  $i$ , i.e.,  $i$  is not  
 802 processing any previously admitted job  $j'$ , the shortest available job  $j^*$  is immediately  
 803 admitted to machine  $i$  at time  $a_j^* := \tau$ . There are two events that trigger a decision of  
 804 the region algorithm: the release of a job and the completion of a job. If one of these  
 805 events occurs at time  $\tau$ , the region algorithm invokes the preemption subroutine. This  
 806 routine iterates over all machines and compares the processing time of the smallest  
 807 job  $j^*$  *available* for machine  $i$  with the processing time of job  $j_i$  that is currently  
 808 scheduled on machine  $i$ . If  $p_{ij^*} < \frac{\varepsilon}{4}p_{ij_i}$ , job  $j^*$  is admitted to machine  $i$  at time  $a_j^* := \tau$   
 809 and, by the above scheduling routine, immediately starts processing. We summarize  
 810 the region algorithm below.

811 The proof of the analysis splits again naturally into two parts: The first part is  
 812 to show that the region algorithm completes at least half of all admitted jobs, and  
 813 the second is to use [Theorem 4.4](#) to compare the number of admitted jobs to the  
 814 throughput of an optimal non-migratory algorithm.

815 **5.2. Completing sufficiently many admitted jobs.** The main result of this  
 816 section is the following theorem.

817 **THEOREM 5.1.** *Let  $0 < \varepsilon \leq 1$ . Then the region algorithm completes at least half*  
 818 *of all admitted jobs before their deadline.*

819 The proof of [Theorem 5.1](#) relies on two technical results that enable us to restrict  
 820 to instances with one machine and further only consider jobs that are admitted by  
 821 the region algorithm in this instance. Then, we can use the analysis of the region  
 822 algorithm in [10] to complete the proof.

823 We start with the following observation. Let  $\mathcal{I}$  be an instance of online throughput  
 824 maximization with the job set  $\mathcal{J}$  and let  $J \subseteq \mathcal{J}$  be the set of jobs admitted by the  
 825 region algorithm at some point. It is easy to see that a job  $j \notin J$  does not influence the  
 826 scheduling or admission decisions of the region algorithm. The next lemma formalizes  
 827 this statement and follows immediately from the just made observations.

828 **LEMMA 5.2.** *For any instance  $\mathcal{I}$  for which the region algorithm admits the job*

**Algorithm** Region algorithm

**Scheduling Routine:** At any time  $\tau$  and on any machine  $i$ , run the job with shortest processing time that has been admitted to  $i$  and has not yet completed.

**Event:** Release of a new job at time  $\tau$   
Call Threshold Preemption Routine.

**Event:** Completion of a job at time  $\tau$   
Call threshold preemption routine.

**Threshold Preemption Routine:**

$i \leftarrow 1$

$j^* \leftarrow$  a shortest job available for machine  $i$  at  $\tau$ , i.e.,  $j^* \in \arg \min\{p_{ij} \mid j \in \mathcal{J}, r_j \leq \tau \text{ and } d_j - \tau \geq (1 + \frac{\varepsilon}{2})p_{ij}\}$

**while**  $i \leq m$  **do**

$j \leftarrow$  job processed on machine  $i$  at time  $\tau$

**if**  $j = \emptyset$  **then**

        admit job  $j^*$  to machine  $i$

        call Threshold Preemption Routine

**else if**  $p_{ij^*} < \frac{\varepsilon}{4}p_{ij}$  **then**

        admit job  $j^*$  to machine  $i$

        call Threshold Preemption Routine

**else**

$i \leftarrow i + 1$

$j^* \leftarrow$  a shortest job available for machine  $i$  at  $\tau$ , i.e.,  $j^* \in \arg \min\{p_{ij} \mid j \in \mathcal{J}, r_j \leq \tau \text{ and } d_j - \tau \geq (1 + \frac{\varepsilon}{2})p_{ij}\}$

**end if**

**end while**

829 set  $J \subseteq \mathcal{J}$ , the reduced instance  $\mathcal{I}'$  containing only the jobs  $J$  forces the region al-  
830 gorithm with consistent tie breaking to admit all jobs in  $J$  and to create the same  
831 schedule as produced for the instance  $\mathcal{I}$ .

832 The proof of the main result compares the number of jobs finished on time,  $F \subseteq J$ ,  
833 to the number of jobs unfinished by their respective deadlines,  $U = J \setminus F$ . To further  
834 simplify the instance, we use that the region algorithm is non-migratory and restrict  
835 to single-machine instances. To this end, let  $F^{(i)}$  and  $U^{(i)}$  denote the finished and  
836 unfinished, respectively, jobs on machine  $i$ .

837 **LEMMA 5.3.** *Let  $i \in \{1, \dots, m\}$ . There is an instance  $\mathcal{I}'$  on one machine with job*  
838 *set  $\mathcal{J}' = F^{(i)} \cup U^{(i)}$ . Moreover, the schedule of the region algorithm for instance  $\mathcal{I}'$*   
839 *with consistent tie breaking is identical to the schedule of the jobs  $\mathcal{J}'$  on machine  $i$ .*  
840 *In particular,  $F' = F^{(i)}$  and  $U' = U^{(i)}$ .*

841 *Proof.* By [Lemma 5.2](#), we can restrict to the jobs admitted by the region algo-  
842 rithm. Hence, let  $\mathcal{I}$  be such an instance with  $F^{(i)} \cup U^{(i)}$  being admitted to machine  $i$ .  
843 As the region algorithm is non-migratory, the sets of jobs scheduled on two different  
844 machines are disjoint. Let  $\mathcal{I}'$  consist of the jobs in  $\mathcal{J}' := F^{(i)} \cup U^{(i)}$  and one machine.  
845 We set  $p'_j = p_{ij}$  for  $j \in \mathcal{J}'$ . The region algorithm on instance  $\mathcal{I}$  admits all jobs in  $\mathcal{J}$ .  
846 In particular, it admits all jobs in  $\mathcal{J}'$  to machine  $i$ .

847 We inductively show that the schedule for the instance  $\mathcal{I}'$  is identical to the  
848 schedule on machine  $i$  in instance  $\mathcal{I}$ . To this end, we index the jobs in  $\mathcal{J}'$  in increasing  
849 admission time points in instance  $\mathcal{I}$ .

850 It is obvious that job  $1 \in \mathcal{J}'$  is admitted to the single machine at its release date  $r_1$   
 851 as happens in instance  $\mathcal{I}$  since the region algorithm uses consistent tie breaking.  
 852 Suppose that the schedule is identical until the admission of job  $j^*$  at time  $a_j^* = \tau$ .  
 853 If  $j^*$  does not interrupt the processing of another job, then  $j^*$  will be admitted at  
 854 time  $\tau$  in  $\mathcal{I}'$  as well. Otherwise, let  $j \in \mathcal{J}'$  be the job that the region algorithm  
 855 planned to process at time  $\tau$  before the admission of job  $j^*$ . Since  $j^*$  is admitted at  
 856 time  $\tau$  in  $\mathcal{I}$ ,  $j^*$  is available at time  $\tau$ , satisfies  $p'_{j^*} = p_{ij^*} < \frac{\varepsilon}{4}p_{ij} = \frac{\varepsilon}{4}p'_j$ , and did not  
 857 satisfy both conditions at some earlier time  $\tau'$  with some earlier admitted job  $j'$ . Since  
 858 the job set in  $\mathcal{I}'$  is a subset of the jobs in  $\mathcal{I}$  and we use consistent tie breaking, no other  
 859 job  $j^* \in \mathcal{J}'$  that satisfies both conditions is favored by the region algorithm over  $j^*$ .  
 860 Therefore, job  $j^*$  is also admitted at time  $\tau$  by the region algorithm in instance  $\mathcal{I}'$ .  
 861 Thus, the schedule created by the region algorithm for  $\mathcal{J}'$  is identical to the schedule  
 862 of  $\mathcal{J}$  on machine  $i$  in the original instance.  $\square$

863 For proving [Theorem 5.1](#), we consider a worst-case instance for the region algo-  
 864 rithm where “worst” is with respect to the ratio between admitted and successfully  
 865 completed jobs. Since the region algorithm is non-migratory, there exists at least one  
 866 machine in such a worst-case instance that “achieves” the same ratio as the whole  
 867 instance. By the just proven lemma, we can find a worst-case instance on a single  
 868 machine. However, on a single machine, the region algorithm algorithm in this paper  
 869 is identical to the algorithm designed in [\[10\]](#). Therefore, we simply follow the line of  
 870 proof developed in [\[10\]](#) to show [Theorem 5.1](#).

871 More precisely, in [\[10\]](#) we show that the existence of a late job  $j$  implies that the  
 872 the set of jobs admitted by  $j$  or by one of its children contains more finished than  
 873 unfinished jobs. Let  $F_j$  denote the set of jobs admitted by  $j$  or by one of its children  
 874 that *finish on time*. Similarly, we denote the set of such jobs that complete after their  
 875 deadlines, i.e., that are *unfinished at their deadline*, by  $U_j$ . We restate the following  
 876 lemma, which was originally shown in a single-machine environment but clearly also  
 877 holds for unrelated machines.

878 **LEMMA 5.4** (Lemma 3 in [\[10\]](#)). *Consider some job  $j$  admitted to some machine*  
 879  *$i \in \{1, \dots, m\}$ . If  $C_j - a_j \geq (\ell + 1)p_{ij}$  for  $\ell > 0$ , then  $|F_j| - |U_j| \geq \lfloor \frac{4\ell}{\varepsilon} \rfloor$ .*

880 *Proof of [Theorem 5.1](#).* Let  $U$  be the set of jobs that are unfinished by their dead-  
 881 line but whose ancestors have all completed on time. Every job  $j \in U$  was admitted  
 882 by the algorithm at some time  $a_j$  with  $d_j - a_j \geq (1 + \frac{\varepsilon}{2})p_{ij}$ . Since  $j$  is unfinished, we  
 883 have  $C_j - a_j > d_j - a_j \geq (1 + \frac{\varepsilon}{2})p_{ij}$ . By [Lemma 5.4](#),  $|F_j| - |U_j| \geq \lfloor \frac{4 \cdot \varepsilon/2}{\varepsilon} \rfloor = 2$ . Thus,

$$884 \quad |F_j| + |U_j| \leq 2|F_j| - 2 < 2|F_j|.$$

885 Since every ancestor of such a job  $j$  finishes on time, this completes the proof.  $\square$

886 **5.3. The region algorithm admits sufficiently many jobs.** In this section,  
 887 we show the following theorem and give the proof of [Theorem 1.4](#).

888 **THEOREM 5.5.** *An optimal non-migratory (offline) algorithm completes at most*  
 889 *a factor  $(\frac{8}{\varepsilon} + 4)$  more jobs on time than admitted by the region algorithm.*

890 *Proof.* As in the previous section, fix an instance and an optimal solution OPT.  
 891 Let  $X$  be the set of jobs in OPT that the region algorithm did not admit. We assume  
 892 without loss of generality that all jobs in OPT finish on time. Further, let  $J$  denote  
 893 the set of jobs that the region algorithm admitted. Then,  $X \cup J$  is a superset of the  
 894 jobs in OPT. Thus,  $|X| \leq (\frac{8}{\varepsilon} + 3)|J|$  implies [Theorem 5.5](#).

895 Consider an arbitrary but fixed machine  $i$ . We compare again the throughput  
 896 of the optimal schedule on machine  $i$  to the throughput of the region algorithm on  
 897 machine  $i$ . Let  $X_i \subseteq X$  denote the jobs in OPT scheduled on machine  $i$  and let  $J_i$   
 898 denote the jobs scheduled by the region algorithm on machine  $i$ . Then, showing  $|X_i| \leq$   
 899  $(\frac{8}{\varepsilon} + 3)|J_i|$  suffices to prove the main result of this section. Given that the region  
 900 algorithm satisfies Properties (P1) and (P2), Theorem 4.2 already provides a bound  
 901 on the cardinality of  $X_i$  in terms of the *intervals* corresponding to the schedule on  
 902 machine  $i$ . Thus, it remains to show that the region algorithm indeed qualifies for  
 903 applying Theorem 4.2 and that the bound developed therein can be translated to a  
 904 bound in terms of  $|J_i|$ .

905 We start by showing that the region algorithm satisfies the assumptions necessary  
 906 for applying Theorem 4.2. Clearly, as the region algorithm only admits a job  $j$   
 907 at time  $\tau$  if  $d_j - \tau \geq (1 + \frac{\varepsilon}{2})p_{ij}$ , setting  $\delta = \frac{\varepsilon}{2}$  proves that the region algorithm  
 908 satisfies (P1). For (P2), we retrospectively analyze the schedule generated by the  
 909 region algorithm. For a time  $\tau$ , let  $j_i$  denote the job scheduled on machine  $i$ . Then,  
 910 setting  $u_{i,\tau} := \frac{\varepsilon}{4}p_{ij_i}$  or  $u_{i,\tau} = \infty$  if no such job  $j_i$  exists, indeed provides us with the  
 911 machine-dependent threshold necessary for (P2). This discussion also implies that  $u^{(i)}$   
 912 has only countably many points of discontinuity as there are only finitely many jobs  
 913 in the instance, and that  $u^{(i)}$  is right-continuous.

914 Hence, let  $\mathcal{I}$  denote the set of maximal intervals  $I_t = [\tau_t, \tau_{t+1})$  for  $t \in [T]$  of  
 915 constant threshold  $u_{i\tau}$ . Thus, by Theorem 4.2,

$$916 \quad (5.1) \quad |X_i| \leq \sum_{t=1}^T \frac{\varepsilon}{\varepsilon - \delta} \frac{\tau_{t+1} - \tau_t}{u_t} + T.$$

917 As the threshold  $u_{i,\tau}$  is proportional to the processing time of the job currently  
 918 scheduled on machine  $i$ , the interval  $I_t$  either represents an idle interval of machine  $i$   
 919 (with  $u_{i\tau} = \infty$ ) or corresponds to the uninterrupted processing of some job  $j$  on  
 920 machine  $i$ . We denote this job by  $j_t$  if it exists. We consider now the set  $\mathcal{I}_j \subseteq \mathcal{I}$   
 921 of intervals with  $j_t = j$  for some particular job  $j \in J_i$ . As observed, these intervals  
 922 correspond to job  $j$  being processed which happens for a total of  $p_{ij}$  units of time.  
 923 Combining with  $u_t = \frac{\varepsilon}{4}p_{ij}$  for  $I_t \in \mathcal{I}_j$ , we get

$$924 \quad \sum_{t: I_t \in \mathcal{I}_j} \frac{\tau_{t+1} - \tau_t}{u_t} = \frac{p_{ij}}{\frac{\varepsilon}{4}p_{ij}} = \frac{4}{\varepsilon}.$$

925 As  $\delta = \frac{\varepsilon}{2}$ , we additionally have that  $\frac{\varepsilon}{\varepsilon - \delta} = 2$ . Hence, we rewrite Equation (5.1) by

$$926 \quad |X_i| \leq \frac{8}{\varepsilon}|J_i| + T.$$

927 It remains to bound  $T$  in terms of  $|J_i|$  to conclude the proof. To this end, we  
 928 recall that the admission of a job  $j$  to a machine interrupts the processing of at most  
 929 one previously admitted job. Hence, the admission of  $|J_i|$  jobs to machine 1 creates  
 930 at most  $2|J_i| + 1$  intervals.

931 If the region algorithm does not admit any job to machine  $i$ , i.e.,  $|J_i| = 0$ ,  
 932 then  $u_{i\tau} = \infty$  for each time point  $\tau$ . Hence, there exists no job scheduled on ma-  
 933 chine  $i$  by OPT that the region algorithm did not admit. In other words,  $X_i = \emptyset$   
 934 and  $|X_i| = 0 = |J_i|$ . Otherwise,  $2|J_i| + 1 \leq 3|J_i|$ . Therefore,

$$935 \quad |X_i| \leq \left(\frac{8}{\varepsilon} + 3\right)|J_i|.$$

936 Combining with the observation about  $X_i$  and  $J_i$  previously discussed, we obtain

$$937 \quad |\text{OPT}| \leq |X \cup J| = \sum_{i=1}^m |X_i| + |J| \leq \left(\frac{8}{\varepsilon} + 3\right) \sum_{i=1}^m |J_i| + |J| = \left(\frac{8}{\varepsilon} + 4\right) |J|,$$

938 which concludes the proof.  $\square$

#### 939 5.4. Finalizing the proof of [Theorem 1.4](#).

940 *Proof of [Theorem 1.4](#).* In [Theorem 5.1](#) we show that the region algorithm com-  
 941 pletes at least half of all admitted jobs  $J$  on time. In [Theorem 4.2](#), we bound  
 942 the throughput  $|\text{OPT}|$  of an optimal non-migratory solution by  $(\frac{8}{\varepsilon} + 4)|J|$ . Com-  
 943 bining these theorems shows that the region algorithm achieves a competitive ratio  
 944 of  $c = 2 \cdot (\frac{8}{\varepsilon} + 4) = \frac{16}{\varepsilon} + 8$ .  $\square$

945 **6. Conclusion.** In this paper, we close the problem of online single-machine  
 946 throughput maximization with and without commitment requirements. For both com-  
 947 mitment settings, we give an optimal online algorithm. Further, our algorithms run  
 948 in a multiple-machine environment, even on heterogenous machines. Our algorithms  
 949 compute non-migratory schedules on unrelated machines with the same competitive  
 950 ratio  $\mathcal{O}(\frac{1}{\varepsilon})$  as for a single machine and improve substantially upon the state of the  
 951 art.

952 It remains open whether the problem with a large number of machines admits an  
 953 online algorithm with a better competitive ratio. For  $m \geq 2$ , it is not known whether  
 954 slack is actually needed to design algorithms with bounded competitive ratios, even  
 955 without commitment requirements and identical machines. In fact, results in [\[26\]](#)  
 956 (used to show a  $\mathcal{O}(1)$ -competitive randomized algorithm on a single machine) imply  
 957 an  $\mathcal{O}(1)$ -competitive algorithm for scheduling jobs without slack and without commit-  
 958 ment on  $m \in \mathcal{O}(1)$  identical machines. Further, for machine utilization, i.e., weighted  
 959 throughput with  $p_j = w_j$ , [\[23, 34\]](#) improve upon the factor of  $\mathcal{O}(\frac{1}{\varepsilon})$  for commitment  
 960 upon arrival and jobs satisfying the  $\varepsilon$ -slack assumption.

961 In fact, there are examples in the literature in which the worst-case ratio for  
 962 a scheduling problem improves with an increasing number of machines. Consider,  
 963 e.g., the non-preemptive offline variant of our throughput maximization problem on  
 964 identical machines. There is an algorithm with approximation ratio of 1.55 for any  $m$   
 965 which is improving with increasing number of machines, converging to 1 as  $m$  tends  
 966 to infinity [\[20\]](#). The second part of the result also holds for the weighted problem.

967 Another interesting question asks whether randomization allows for improved re-  
 968 sults. Recall that there is a  $\mathcal{O}(1)$ -competitive randomized algorithm for scheduling on  
 969 a single machine without commitment and without slack assumption [\[26\]](#). Therefore  
 970 it seems plausible that randomization also helps designing algorithms with improved  
 971 competitive ratios for the different commitment models, for which only weak lower  
 972 bounds are known [\[10\]](#), and on multiple machines as discussed above.

973 Further, we leave migratory scheduling on unrelated machines as an open prob-  
 974 lem. Allowing migration in this setting means that, on each machine  $i$ , a certain  
 975 fraction of the processing time  $p_{i,j}$  is executed, and these fractions must sum to one.  
 976 Generalizing the result we leverage for identical machines [\[25\]](#), it is conceivable that  
 977 any migratory schedule can be turned into a valid non-migratory schedule of the same  
 978 jobs by adding a constant number of machines *of each type*. Such a result would im-  
 979 mediately allow to transfer our competitive ratios to the migratory setting (up to  
 980 constant factors). Devanur and Kulkarni [\[13\]](#) show a weaker result that utilizes speed  
 981 rather than additional machines. Note that the strong impossibility result of Im and



982 Moseley [22] does not rule out the desired strengthening because we make the  $\varepsilon$ -slack  
 983 assumption for every job and machine eligible for it. Further, we – as well as Devanur  
 984 and Kulkarni [13] – assume that the processing time of each job  $j$  satisfies  $p_{ij} \leq d_j - r_j$   
 985 on any eligible machine  $i$ , whereas the lower bound in [22] requires jobs that violate  
 986 this reasonable assumption.

987 Further research directions include generalizations such as weighted throughput  
 988 maximization. While strong lower bounds exist for handling weighted throughput  
 989 with commitment [10], there remains a gap for the problem without. The known  
 990 lower bound of  $\Omega(\frac{1}{\varepsilon})$  already holds for unit weights [10]. A natural extension of  
 991 the region algorithm bases its admission decisions on the density, i.e., the ratio of the  
 992 weight of a job to its processing time. The result is an algorithm similar to the  $\mathcal{O}(\frac{1}{\varepsilon^2})$ -  
 993 competitive algorithm by Lucier et al. [31]. Both algorithms only admit available jobs  
 994 and interrupt currently running jobs if the new job is denser by a certain factor.  
 995 However, we can show that there is a lower bound of  $\Omega(\frac{1}{\varepsilon^2})$  on the competitive ratio  
 996 of such algorithms. Hence, in order to improve the upper bound for online weighted  
 997 throughput maximization, one needs to develop a new type of algorithm.

998

## REFERENCES

- 999 [1] K. AGRAWAL, J. LI, K. LU, AND B. MOSELEY, *Scheduling parallelizable jobs online to maximize*  
 1000 *throughput*, in LATIN, vol. 10807, Springer, 2018, pp. 755–776, [https://doi.org/10.1007/](https://doi.org/10.1007/978-3-319-77404-6_55)  
 1001 [978-3-319-77404-6\\_55](https://doi.org/10.1007/978-3-319-77404-6_55).
- 1002 [2] Y. AZAR, I. KALP-SHALTIEL, B. LUCIER, I. MENACHE, J. NAOR, AND J. YANIV, *Truthful online*  
 1003 *scheduling with commitments*, in EC, ACM, 2015, pp. 715–732, [https://doi.org/10.1145/](https://doi.org/10.1145/2764468.2764535)  
 1004 [2764468.2764535](https://doi.org/10.1145/2764468.2764535).
- 1005 [3] N. BANSAL, H. CHAN, AND K. PRUHS, *Competitive algorithms for due date scheduling*, Algo-  
 1006 *rithmica*, 59 (2011), pp. 569–582, <https://doi.org/10.1007/s00453-009-9321-4>.
- 1007 [4] A. BAR-NOY, S. GUHA, J. NAOR, AND B. SCHIEBER, *Approximating the throughput of multiple*  
 1008 *machines in real-time scheduling*, SIAM J. Comput., 31 (2001), pp. 331–352, [https://doi.](https://doi.org/10.1137/S0097539799354138)  
 1009 [org/10.1137/S0097539799354138](https://doi.org/10.1137/S0097539799354138), <https://doi.org/10.1137/S0097539799354138>.
- 1010 [5] S. K. BARUAH AND J. R. HARITSA, *Scheduling for overload in real-time systems*, IEEE Trans.  
 1011 *Computers*, 46 (1997), pp. 1034–1039, <https://doi.org/10.1109/12.620484>.
- 1012 [6] S. K. BARUAH, J. R. HARITSA, AND N. SHARMA, *On-line scheduling to maximize task com-*  
 1013 *pletions*, in RTSS, IEEE Computer Society, 1994, pp. 228–236, [https://doi.org/10.1109/](https://doi.org/10.1109/REAL.1994.342713)  
 1014 [REAL.1994.342713](https://doi.org/10.1109/REAL.1994.342713).
- 1015 [7] S. K. BARUAH, G. KOREN, D. MAO, B. MISHRA, A. RAGHUNATHAN, L. E. ROSIER, D. E.  
 1016 *SHASHA*, AND F. WANG, *On the competitiveness of on-line real-time task scheduling*, Real-  
 1017 *Time Systems*, 4 (1992), pp. 125–144, <https://doi.org/10.1007/BF00365406>.
- 1018 [8] S. K. BARUAH, G. KOREN, B. MISHRA, A. RAGHUNATHAN, L. E. ROSIER, AND D. E. SHASHA,  
 1019 *On-line scheduling in the presence of overload*, in FOCS, IEEE Computer Society, 1991,  
 1020 pp. 100–110, <https://doi.org/10.1109/SFCS.1991.185354>.
- 1021 [9] P. BERMAN AND B. DASGUPTA, *Improvements in throughout maximization for real-time sched-*  
 1022 *uling*, in STOC, ACM, 2000, pp. 680–687.
- 1023 [10] L. CHEN, F. EBERLE, N. MEGOW, K. SCHEWIOR, AND C. STEIN, *A general framework for han-*  
 1024 *dling commitment in online throughput maximization*, Math. Prog., 183 (2020), pp. 215–  
 1025 247, <https://doi.org/10.1007/s10107-020-01469-2>.
- 1026 [11] B. DASGUPTA AND M. A. PALIS, *Online real-time preemptive scheduling of jobs with deadlines*,  
 1027 in APPROX, vol. 1913 of Lecture Notes in Computer Science, Springer, 2000, pp. 96–107,  
 1028 [https://doi.org/10.1007/3-540-44436-X\\_11](https://doi.org/10.1007/3-540-44436-X_11).
- 1029 [12] M. L. DERTOUZOS AND A. K. MOK, *Multiprocessor on-line scheduling of hard-real-time tasks*,  
 1030 IEEE Trans. Software Eng., 15 (1989), pp. 1497–1506, <https://doi.org/10.1109/32.58762>.
- 1031 [13] N. R. DEVANUR AND J. KULKARNI, *A unified rounding algorithm for unrelated machines sched-*  
 1032 *uling problems*, in SPAA, C. Scheideler and J. T. Fineman, eds., ACM, 2018, pp. 283–290,  
 1033 <https://doi.org/10.1145/3210377.3210384>, <https://doi.org/10.1145/3210377.3210384>.
- 1034 [14] J. DU AND J. Y. LEUNG, *Minimizing the number of late jobs on unrelated machines*, Oper.  
 1035 *Res. Lett.*, 10 (1991), pp. 153–158.
- 1036 [15] F. EBERLE, N. MEGOW, AND K. SCHEWIOR, *Optimally handling commitment issues in online*

- 1037 *throughput maximization*, in Proceedings of ESA, vol. 173 of LIPIcs, Schloss Dagstuhl -  
1038 Leibniz-Zentrum für Informatik, 2020.
- 1039 [16] A. D. FERGUSON, P. BODÍK, S. KANDULA, E. BOUTIN, AND R. FONSECA, *Jockey: Guaranteed*  
1040 *job latency in data parallel clusters*, in EuroSys, ACM, 2012, pp. 99–112, [https://doi.org/](https://doi.org/10.1145/2168836.2168847)  
1041 [10.1145/2168836.2168847](https://doi.org/10.1145/2168836.2168847).
- 1042 [17] J. A. GARAY, J. NAOR, B. YENER, AND P. ZHAO, *On-line admission control and packet*  
1043 *scheduling with interleaving*, in INFOCOM, IEEE Computer Society, 2002, pp. 94–103,  
1044 <https://doi.org/10.1109/INFCOM.2002.1019250>.
- 1045 [18] M. H. GOLDWASSER, *Patience is a virtue: The effect of slack on competitiveness for admission*  
1046 *control*, J. Sched., 6 (2003), pp. 183–211, <https://doi.org/10.1023/A:1022994010777>.
- 1047 [19] M. H. GOLDWASSER AND B. KERBIKOV, *Admission control with immediate notification*, J.  
1048 Sched., 6 (2003), pp. 269–285, <https://doi.org/10.1023/A:1022956425198>.
- 1049 [20] S. IM, S. LI, AND B. MOSELEY, *Breaking 1 - 1/e barrier for non-preemptive through-*  
1050 *put maximization*, in IPCO, F. Eisenbrand and J. Könemann, eds., vol. 10328 of Lec-  
1051 ture Notes in Computer Science, Springer, 2017, pp. 292–304, [https://doi.org/10.1007/](https://doi.org/10.1007/978-3-319-59250-3_24)  
1052 [978-3-319-59250-3\\_24](https://doi.org/10.1007/978-3-319-59250-3_24), [https://doi.org/10.1007/978-3-319-59250-3\\_24](https://doi.org/10.1007/978-3-319-59250-3_24).
- 1053 [21] S. IM, S. LI, AND B. MOSELEY, *Breaking 1 - 1/e barrier for nonpreemptive throughput max-*  
1054 *imization*, SIAM J. Discret. Math., 34 (2020), pp. 1649–1669, [https://doi.org/10.1137/](https://doi.org/10.1137/17M1148438)  
1055 [17M1148438](https://doi.org/10.1137/17M1148438), <https://doi.org/10.1137/17M1148438>.
- 1056 [22] S. IM AND B. MOSELEY, *General profit scheduling and the power of migration on heterogeneous*  
1057 *machines*, in SPAA, vol. 10807 of Lecture Notes in Computer Science, Springer, 2018,  
1058 pp. 755–776, [https://doi.org/10.1007/978-3-319-77404-6\\_55](https://doi.org/10.1007/978-3-319-77404-6_55).
- 1059 [23] S. JAMALABADI, C. SCHWIEGELSHOHN, AND U. SCHWIEGELSHOHN, *Commitment and slack for*  
1060 *online load maximization*, in SPAA, ACM, 2020, pp. 339–348, [https://doi.org/10.1145/](https://doi.org/10.1145/3350755.3400271)  
1061 [3350755.3400271](https://doi.org/10.1145/3350755.3400271).
- 1062 [24] B. KALYANASUNDARAM AND K. PRUHS, *Speed is as powerful as clairvoyance*, J. ACM, 47 (2000),  
1063 pp. 617–643, <https://doi.org/10.1145/347476.347479>.
- 1064 [25] B. KALYANASUNDARAM AND K. PRUHS, *Eliminating migration in multi-processor scheduling*, J.  
1065 Algorithms, 38 (2001), pp. 2–24, <https://doi.org/10.1006/jagm.2000.1128>.
- 1066 [26] B. KALYANASUNDARAM AND K. PRUHS, *Maximizing job completions online*, J. Algorithms, 49  
1067 (2003), pp. 63–85, [https://doi.org/10.1016/S0196-6774\(03\)00074-9](https://doi.org/10.1016/S0196-6774(03)00074-9).
- 1068 [27] G. KOREN AND D. E. SHASHA, *MOCA: A multiprocessor on-line competitive algorithm for*  
1069 *real-time system scheduling*, Theor. Comput. Sci., 128 (1994), pp. 75–97, [https://doi.org/](https://doi.org/10.1016/0304-3975(94)90165-1)  
1070 [10.1016/0304-3975\(94\)90165-1](https://doi.org/10.1016/0304-3975(94)90165-1).
- 1071 [28] G. KOREN AND D. E. SHASHA, *D<sup>over</sup>: An optimal on-line scheduling algorithm for overloaded*  
1072 *uniprocessor real-time systems*, SIAM J. Comput., 24 (1995), pp. 318–339, [https://doi.](https://doi.org/10.1137/S0097539792236882)  
1073 [org/10.1137/S0097539792236882](https://doi.org/10.1137/S0097539792236882).
- 1074 [29] E. LAWLER, *A dynamic programming algorithm for preemptive scheduling of a single machine*  
1075 *to minimize the number of late jobs*, Ann. Oper. Res., 26 (1990), pp. 125–133.
- 1076 [30] E. L. LAWLER, *Recent results in the theory of machine scheduling*, in Mathematical Pro-  
1077 gramming The State of the Art, A. Bachem, B. Korte, and M. Grötschel, eds., Springer,  
1078 1982, pp. 202–234, [https://doi.org/10.1007/978-3-642-68874-4\\_9](https://doi.org/10.1007/978-3-642-68874-4_9), [https://doi.org/10.1007/](https://doi.org/10.1007/978-3-642-68874-4_9)  
1079 [978-3-642-68874-4\\_9](https://doi.org/10.1007/978-3-642-68874-4_9).
- 1080 [31] B. LUCIER, I. MENACHE, J. NAOR, AND J. YANIV, *Efficient online scheduling for deadline-*  
1081 *sensitive jobs: Extended abstract*, in SPAA, ACM, 2013, pp. 305–314, [https://doi.org/10.](https://doi.org/10.1145/2486159.2486187)  
1082 [1145/2486159.2486187](https://doi.org/10.1145/2486159.2486187).
- 1083 [32] K. PRUHS AND C. STEIN, *How to schedule when you have to buy your energy*, in APPROX,  
1084 vol. 6302 of Lecture Notes in Computer Science, Springer, 2010, pp. 352–365, [https://doi.](https://doi.org/10.1007/978-3-642-15369-3_27)  
1085 [org/10.1007/978-3-642-15369-3\\_27](https://doi.org/10.1007/978-3-642-15369-3_27).
- 1086 [33] K. PRUHS AND G. J. WOEGINGER, *Approximation schemes for a class of subset selection prob-*  
1087 *lems*, Theor. Comput. Sci., 382 (2007), pp. 151–156, <https://doi.org/10.1016/j.tcs.2007.03.006>,  
1088 <https://doi.org/10.1016/j.tcs.2007.03.006>.
- 1089 [34] C. SCHWIEGELSHOHN AND U. SCHWIEGELSHOHN, *The power of migration for online slack sched-*  
1090 *uling*, in ESA, vol. 57 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016,  
1091 pp. 75:1–75:17, <https://doi.org/10.4230/LIPIcs.ESA.2016.75>.
- 1092 [35] R. SITTERS, *Complexity of preemptive minsum scheduling on unrelated parallel machines*, Jour-  
1093 nal of Algorithms, 57 (2005), pp. 37–48, [https://doi.org/https://doi.org/10.1016/j.jalgor.](https://doi.org/https://doi.org/10.1016/j.jalgor.2004.06.011)  
1094 [2004.06.011](https://doi.org/https://doi.org/10.1016/j.jalgor.2004.06.011), <https://www.sciencedirect.com/science/article/pii/S0196677404001130>.