

# Online Scheduling with Bounded Migration

Peter Sanders

Universität Karlsruhe (TH), Fakultät für Informatik, Postfach 6980, 76128 Karlsruhe, Germany  
email: sanders@ira.uka.de <http://algo2.iti.uni-karlsruhe.de/sanders.php>

Naveen Sivadasan

Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany  
email: ns@mpi-sb.mpg.de <http://www.mpi-inf.mpg.de/~sivadasa/>

Martin Skutella

TU Berlin, Institut für Mathematik, Straße des 17. Juni 136, 10623 Berlin, Germany  
email: skutella@math.tu-berlin.de <http://www.math.tu-berlin.de/~skutella/>

Consider the classical online scheduling problem where jobs that arrive one by one are assigned to identical parallel machines with the objective of minimizing the makespan. We generalize this problem by allowing the current assignment to be changed whenever a new job arrives, subject to the constraint that the total size of moved jobs is bounded by  $\beta$  times the size of the arriving job. For small values of  $\beta$ , we obtain several simple online algorithms with constant competitive ratio. We also present a linear time ‘online approximation scheme’, that is, a family of online algorithms with competitive ratio  $1 + \epsilon$  and constant migration factor  $\beta(\epsilon)$ , for any fixed  $\epsilon > 0$ .

*Key words:* scheduling; approximation; online algorithm; sensitivity analysis

*MSC2000 Subject Classification:* Primary: 90C27, 90C31, 68Q25; Secondary: 90C10, 68M20

*OR/MS subject classification:* Primary: production/scheduling, sequencing, multiple machine; Secondary: integer programming

---

**1. Introduction** One of the most fundamental scheduling problems asks for an assignment of jobs to  $m$  identical parallel machines so as to minimize the makespan. (The makespan is the completion time of the last job that finishes in the schedule; it also equals the maximum machine load.) In the standard classification scheme of Graham, Lawler, Lenstra, & Rinnooy Kan [16], this scheduling problem is denoted by  $P \parallel C_{\max}$  and it is well known to be strongly NP-hard [13].

The *offline* variant of this problem assumes that all jobs are known in advance whereas in the *online* variant the jobs are incrementally revealed by an adversary and the online algorithm can only choose the machine for the new job without being allowed to move other jobs. Note that dropping this radical constraint on the online algorithm yields the offline situation.

We study a natural generalization of both offline and online problems. Jobs arrive incrementally but, upon arrival of a new job  $j$ , we are allowed to migrate *some* previous jobs to other machines. The total size of the migrated jobs however must be bounded by  $\beta p_j$  where  $p_j$  is the size of the new job. For *migration factor*  $\beta = 0$  we get the online setting and for  $\beta = \infty$  we get the offline setting.

For an offline optimization problem, an *approximation algorithm* efficiently (in polynomial time) constructs schedules whose values are within a constant factor  $\alpha$  of the optimum solution value. The number  $\alpha$  is called *performance guarantee* or *performance ratio* of the approximation algorithm. A family of polynomial time approximation algorithms with performance guarantee  $1 + \epsilon$  for all fixed  $\epsilon > 0$  is called a *polynomial time approximation scheme* (PTAS).

In a similar way, *competitive analysis* evaluates solutions computed in the online setting. An online algorithm achieves *competitive ratio*  $\alpha$  if it always maintains solutions whose objective values are within a factor  $\alpha$  of the offline optimum. Here, in contrast to offline approximation results, the achievable values  $\alpha$  are not determined by limited computing power but by the apparent lack of information about parts of the input that will only be revealed in the future. As a consequence, for all interesting classical online problems, it is rather easy to come up with lower bounds that create a gap between the best possible competitive ratio  $\alpha$  and 1. In particular, it is usually impossible to construct a family of  $(1+\epsilon)$ -competitive online algorithms for such problems.

**Related Work** For the online machine scheduling problem, Graham’s *list scheduling* algorithm keeps the makespan within a factor  $2 - \frac{1}{m}$  of the offline optimum [14]: Schedule a newly arriving job on a least loaded machine. It can also easily be seen that this bound is tight. For the offline setting, Graham [15] shows that sorting the jobs in the order of non-increasing size before feeding them to the list scheduling algorithm yields an approximation algorithm with performance ratio  $\frac{4}{3} - \frac{1}{3m}$ . Later, exploiting the relationship between the machine scheduling problem under consideration and the bin-packing problem, algorithms with improved approximation ratios have been obtained in a series of works [9, 12, 20].

Finally, PTASes for a constant number of machines and for an arbitrary number of machines are given in [15, 25] and by Hochbaum and Shmoys [18], respectively. The latter PTAS partitions jobs into large and small jobs. The sizes of large jobs are rounded such that an optimum schedule for the rounded jobs can be obtained via dynamic programming. The small jobs are then added greedily using Graham’s list scheduling algorithm. This approach can be refined to an algorithm with linear running time (see, e.g., [17]): replace the dynamic program with an integer linear program on a fixed number of variables and constraints which can be solved in constant time [21].

In a series of papers, increasingly complicated online algorithms with better and better competitive ratios beating the Graham bound 2 have been developed [6, 19, 2]. The best result known to date is a 1.9201-competitive algorithm due to Fleischer and Wahl [11]. The best lower bound 1.88 on the competitive ratio of any deterministic online algorithm currently known is due to Rudin and Chandrasekaran [24] (see also Rudin [23]). For randomized online algorithms there is a lower bound of  $\frac{e}{e-1} \approx 1.58$  [8, 29]. For more results on online algorithms for scheduling we refer to the recent survey articles by Albers [3] and Sgall [30].

Strategies that reassign jobs were studied in the context of online load balancing where jobs arrive in and depart from a system of  $m$  machines online and the scheduler has to assign each incoming job to one of the machines. Deviating from the usual approach of comparing against the optimal *peak load* seen so far, Westbrook [31] introduced the notion of competitiveness against *current load*: An algorithm is  $\alpha$ -competitive if after every round the makespan is within  $\alpha$  factor of the optimal makespan for the current set of jobs. Each incoming job  $j$  has size  $p_j$  and reassignment cost  $c_j$ . For a job, the reassignment cost has to be paid for its initial assignment and then every time it is reassigned. Observe that the optimal strategy has to pay this cost once for each job for its initial assignment. Thus the optimal (re)assignment cost  $S$  is simply the sum of reassignment costs of all jobs scheduled till now. Westbrook showed a 6-competitive strategy for identical machines with reassignment cost  $3S$  for proportional reassignments, i.e.,  $c_j$  is proportional to  $p_j$ , and  $2S$  for unit reassignments, i.e.,  $c_j = 1$  for all jobs. Later Andrews, Goemans, and Zhang [4] improved it to competitive ratio 3.5981 with the same reassignment factors. They also showed  $(3+\epsilon)$ - and  $(2+\epsilon)$ -competitive strategies for the proportional and unit case, respectively, where the reassignment factor depends only on  $\epsilon$ . For arbitrary reassignment costs they achieve 3.5981-competitiveness with 6.8285 reassignment factor. They also present a 32-competitive strategy with constant reassignment factor for related machines.

**Our Contribution** We consider the online setting where jobs arrive incrementally but there are no job departures. In Section 3 we describe a simple online algorithm which achieves approximation ratio  $\frac{3}{2}$  using a moderate migration factor  $\beta = 2$ . Notice that already this result beats the lower bound 1.88 (1.58) on the competitive ratio of any classical (randomized) online algorithm without migration. Using a more sophisticated analysis, the migration factor can be decreased to  $\frac{4}{3}$  while maintaining competitive ratio  $\frac{3}{2}$ . This result is tight since we can show that any robust<sup>1</sup> scheduling strategy with competitive ratio  $\frac{3}{2}$  has migration factor at least  $\frac{4}{3}$ . We also show another tightness result: The algorithms considered in Section 3 have the nice property that they work *locally*, i.e., they only migrate jobs from the machine where the new job is being assigned to. We prove that any robust local strategy has competitive ratio at least  $\frac{3}{2}$ . In Section 4 we present a more sophisticated algorithm with migration factor 4 and competitive ratio  $\frac{4}{3}$ . For the special case of two machines it is shown in Section 5 that we can achieve competitive ratio  $\frac{4}{3}$  with a migration factor of 1. Moreover, this ratio is tight for migration factor 1.

In Section 6 we present a family of online algorithms with competitive ratio  $1 + \epsilon$  and constant migration factor  $\beta(\epsilon)$ , for any fixed  $\epsilon > 0$ . On the negative side, no constant migration factor suffices to maintain competitive ratio 1, i.e., optimality. We provide interpretations of these results in several

<sup>1</sup>We give a precise definition of robustness in Section 3 after the proof of Theorem 3.2.

different contexts:

- (i) *Online algorithms.* Online scheduling with bounded job migration is a relaxation of the classical online paradigm. Obviously, there is a tradeoff between the desire for high quality solutions and the requirement to compute them online, that is, to deal with a lack of information. Our result can be interpreted in terms of the corresponding tradeoff curve: Any desired quality can be guaranteed while relaxing the online paradigm only moderately by allowing for a constant migration factor.
- (ii) *Sensitivity analysis.* Given an optimum solution to an instance of an optimization problem and a slightly modified instance, can the given solution be turned into an optimum solution for the modified instance without changing the solution too much? This is the compelling question in sensitivity analysis. As indicated above, for the scheduling problem under consideration one has to answer in the negative. Already one additional job can change the entire structure of an optimum schedule. However, our result implies that the answer is positive if we only require near-optimum solutions.
- (iii) *Approximation results.* Our result yields a new PTAS for the scheduling problem under consideration. Due to its online background, this PTAS constructs the solution incrementally. That is, it reads the input little by little always maintaining a  $(1 + \epsilon)$ -approximate solution. Indeed, it follows from the analysis of the algorithm that every update only takes constant time. In particular, the overall running time is linear and thus matches the previously best known approximation result.

We believe that each of these interpretations constitutes an interesting motivation for results like the one we present here in its own right and can therefore lead to interesting results for many other optimization problems. After reading a preliminary version of this paper [27], Epstein and Levin [10] came up with similar results for the bin packing problem.

Our results can also be interpreted within the framework of online load balancing (see the discussion above and [31, 4]), with proportional reassignments and without job deletions. But our requirements are stronger in the sense that a strategy with reassignment factor  $\beta$  ensures that when a new job  $j$  arrives, the total reassignment cost incurred (for scheduling it) is at most  $\beta c_j$ . This is different from the more relaxed constraint that after  $t$  rounds, the total reassignment cost incurred is at most  $\beta \sum c_j$  (summing over all jobs seen till round  $t$ ). Moreover, almost all our scheduling strategies are *robust*, i.e., they convert *any*  $\alpha$ -competitive schedule to an  $\alpha$ -competitive schedule after assigning the newly arrived job, whereas in [31, 4] it is required that the schedule so far is carefully constructed in order to ensure the competitiveness after assigning/deleting a job in the next round.

The underlying details of the presented online approximation scheme have the same roots as the original PTAS by Hochbaum and Shmoys [18] and its refinements [17]. We distinguish between small and large jobs; a job is called large if its size is of the same order of magnitude as the optimum makespan. Since this optimum can change when a new job arrives, the classification of jobs must be updated dynamically. The size of every large job is rounded such that the problem of computing an optimum schedule for the subset of large jobs can be formulated as an integer linear program of constant size. A newly arriving job causes a small change in the right hand side of this program. This enables us to use results from sensitivity analysis of integer programs in order to prove that the schedule of large jobs needs to be changed only slightly. Our PTAS is very simple, it uses only this structural result and does not use any algorithms from integer programming theory.

In Section 7 we discuss an application of bounded migration to configuring storage servers. This was the original motivation for our work. In this application, the objective is to maximize the minimum machine load. It is well-known [5] that any online deterministic algorithm for this *machine covering problem* has competitive ratio at least  $m$  (the number of machines). There is also a lower bound of  $\Omega(\sqrt{m})$  for any randomized online algorithm. We develop a simple deterministic online strategy which is 2-competitive already for migration factor  $\beta = 1$ .

A preliminary version of this work appeared in proceedings of ICALP 2004 [27].

**2. Preliminaries** We consider the problem of scheduling a set of jobs  $\{1, \dots, n\}$  on a set of  $m$  identical parallel machines  $M = \{1, \dots, m\}$ . Job  $j \in \{1, \dots, n\}$  has a positive *processing time* or *size*  $p_j$  and arrives in round  $j$ . The number of jobs is unknown in advance and the processing time of job  $j$  is only revealed in round  $j$ . For a subset of jobs  $N \subseteq \{1, \dots, n\}$ , the *total processing time* of jobs in  $N$  is  $p(N) := \sum_{j \in N} p_j$ ; moreover let  $p_{\max}(N) := \max_{j \in N} p_j$  denote the maximum processing time of a job in  $N$ .

A *schedule*  $S$  for a subset of jobs  $N$  is an assignment of the jobs in  $N$  to the set of machines  $M$ , i.e.,  $S : N \rightarrow M$ . We denote the set of jobs scheduled on (i.e., assigned to) machine  $i \in M$  by  $S^{-1}(i) \subseteq N$ . The *load of machine*  $i$  is  $p(S^{-1}(i))$  and the *makespan* of schedule  $S$  is the maximum machine load  $\max_{i \in M} p(S^{-1}(i))$ . We say that machine  $i \in M$  is a *least loaded* machine with respect to schedule  $S$  if  $p(S^{-1}(i)) = \min_{i' \in M} p(S^{-1}(i'))$ . For a subset of jobs  $N$ , let  $\text{OPT}(N)$  denote the *minimal makespan* of a schedule for  $N$ . The following well-known facts are due to Graham [14].

OBSERVATION 2.1 *Let  $N$  be a set of jobs.*

(i) *The value  $\text{LB}(N) := \max\{\frac{p(N)}{m}, p_{\max}(N)\}$  is a lower bound on  $\text{OPT}(N)$  satisfying*

$$\text{LB}(N) \leq \text{OPT}(N) < 2\text{LB}(N) . \quad (1)$$

(ii) *Consider an arbitrary schedule for  $N$  with makespan  $\kappa$ . Assigning a new job  $j \notin N$  to a least loaded machine yields a schedule for  $N \cup \{j\}$  with makespan at most*

$$\max\{\kappa, \text{OPT}(N \cup \{j\}) + (1 - \frac{1}{m})p_j\} .$$

In the following sections we often consider the subset of jobs  $N = \{1, \dots, j-1\}$  that have arrived in the first  $j-1$  rounds and a newly arrived job  $j$ . If  $N$  and  $j$  are clear from the context, we sometimes use the shorter notation  $\text{OPT} := \text{OPT}(N)$  and  $\text{OPT}' := \text{OPT}(N \cup \{j\})$ .

We conclude this section by stating an easy but important Corollary of Observation 2.1 ((ii)).

COROLLARY 2.1 *Given an  $\alpha$ -approximate schedule for the jobs in  $N$ , assigning a new job of size at most  $(1 + \frac{1}{m-1})(\alpha - 1)\text{OPT}'$  to a least loaded machine yields an  $\alpha$ -approximate schedule for  $N \cup \{j\}$ .*

**3. Simple Strategies with Small Migration Factors** We start by presenting a very simple  $(\frac{3}{2} - \frac{1}{2m})$ -competitive algorithm with migration factor 2. The algorithm is based on the following procedure.

PROCEDURE 1

Input: A schedule for a set of jobs  $N$  and a new job  $j \notin N$ .

Output: A schedule for  $N \cup \{j\}$ .

*Choose one of the following two options which minimizes the resulting makespan:*

Option 0: *Assign job  $j$  to a least loaded machine.*

Option 1: *Let  $i$  be the machine minimizing the maximum job size. Repeatedly remove jobs from this machine (in any order); stop before the total size of removed jobs exceeds  $2p_j$ . Assign job  $j$  to machine  $i$ . Assign the removed jobs successively to a least loaded machine.*

THEOREM 3.1 *Given a  $(\frac{3}{2} - \frac{1}{2m})$ -approximate schedule for  $N$  and a new job  $j \notin N$ , Procedure 1 outputs a  $(\frac{3}{2} - \frac{1}{2m})$ -approximate schedule for  $N \cup \{j\}$ . The migration factor is bounded by 2.*

PROOF. From the description of Procedure 1, it is clear that the migration factor is at most 2. We call a job *small* if its processing time is at most  $\frac{1}{2}\text{OPT}'$ , otherwise it is called *large*. If the new job  $j$  is small, Option 0 yields makespan at most  $(\frac{3}{2} - \frac{1}{2m})\text{OPT}'$  by Corollary 2.1. Thus, we can assume from now on that  $j$  is large.

Since there are at most  $m$  large jobs in  $N \cup \{j\}$ , all jobs on machine  $i$  chosen in Option 1 are small. Thus, after removing jobs from machine  $i$  as described above, machine  $i$  is either empty or the total size of removed jobs exceeds the size of the large job  $j$ . In both cases, assigning job  $j$  to machine  $i$  cannot

increase its load above  $(\frac{3}{2} - \frac{1}{2m})\text{OPT}'$ . Thus, using the same argument as above, assigning the removed small jobs successively to a least loaded machine yields a  $(\frac{3}{2} - \frac{1}{2m})$ -approximate schedule for the set of jobs  $N \cup \{j\}$ .  $\square$

**COROLLARY 3.1** *Scheduling every newly arriving job according to Procedure 1 is a  $(\frac{3}{2} - \frac{1}{2m})$ -competitive algorithm with migration factor 2.*

Next we show that the migration factor can be decreased to  $\frac{4}{3}$  without increasing the competitive ratio above  $\frac{3}{2}$ . This result is achieved by carefully modifying Procedure 1.

**PROCEDURE 2**

Input: A schedule for a set of jobs  $N$  and a new job  $j \notin N$ .

Output: A schedule for  $N \cup \{j\}$ .

Choose one of the following  $m + 1$  options which minimizes the resulting makespan. (Break ties in favor of Option 0.)

Option 0: Assign job  $j$  to a least loaded machine.

Option  $i$  [for  $i \in \{1, \dots, m\}$ ]: Ignoring the largest job on machine  $i$ , consider the remaining jobs on machine  $i$  in order of non-increasing size and remove a job unless the total size of removed jobs will exceed  $\frac{4}{3}p_j$ . Assign job  $j$  to machine  $i$ . Assign the removed jobs successively to a least loaded machine.

In the analysis of Procedure 2 we make use of the following structural property of a schedule. A schedule for a set of jobs  $N$  has property  $(*)$  if the load on any machine excluding its largest job is at most  $\text{OPT}(N)$ . For example, a schedule constructed by list scheduling has always property  $(*)$ . We show that Procedure 2 maintains property  $(*)$ .

**LEMMA 3.1** *Given a schedule for  $N$  satisfying property  $(*)$  and a new job  $j \notin N$ , Procedure 2 outputs a schedule for  $N \cup \{j\}$  satisfying property  $(*)$ .*

**PROOF.** Let  $M' \subseteq M$  be the subset of machines touched by Procedure 2 for scheduling job  $j$ . It suffices to argue that the total load on machine  $i \in M'$  excluding the job  $j'$  that entered the machine last is at most  $\text{OPT}'$ . Except for the very last case considered below, we will prove this stronger property for machine  $i \in M'$ .

If  $j' \neq j$ , then  $j'$  was assigned as part of the redistribution phase. Since redistribution is always performed on a currently least loaded machine, the load on this machine excluding this last job is at most  $\text{OPT}'$ .

It remains to consider the case  $j' = j$ . If Option 0 is chosen and assigns  $j$  to machine  $i$ , then machine  $i$  is a least loaded machine of the initial schedule; in particular, its load excluding  $j$  is at most  $\text{OPT} \leq \text{OPT}'$ . Otherwise,  $j$  is assigned to machine  $i$  by Option  $i$ , which we assume from now on. Since we always break ties in favor of Option 0, the makespan of the schedule produced in Option  $i$  is strictly smaller than the makespan of the schedule obtained in Option 0. We distinguish two cases.

**CASE 1.** The makespan of the schedule obtained by assigning  $j$  to a least loaded machine  $i'$  (Option 0) is determined by machine  $i'$ . As the makespan obtained by Option  $i$  is strictly smaller than the makespan obtained by Option 0, the load of machine  $i$  minus the size of  $j$  is at most the makespan of a least loaded machine  $i'$  in the initial schedule and thus at most  $\text{OPT} \leq \text{OPT}'$ .

**CASE 2.** The makespan of the schedule obtained by assigning  $j$  to a least loaded machine  $i'$  (Option 0) is not determined by machine  $i'$ . In this case, it must be determined by machine  $i \neq i'$  since this is the only machine whose makespan can possibly be decreased by choosing Option  $i$ . Remember that we break ties in favor of Option 0. Thus, the makespan of machine  $i$  must have been decreased in Option  $i$ . Since we did not touch the largest job on machine  $i$ , it still has the desired property.  $\square$

**THEOREM 3.2** *Given a  $\frac{3}{2}$ -approximate schedule for  $N$  satisfying property  $(*)$  and a new job  $j \notin N$ , Procedure 2 outputs a  $\frac{3}{2}$ -approximate schedule for  $N \cup \{j\}$  satisfying property  $(*)$ . The migration factor is bounded by  $\frac{4}{3}$ .*

PROOF. The bound on the migration factor is clear from the description of Procedure 2. Moreover it follows from Lemma 3.1 that the computed schedule for  $N \cup \{j\}$  has property (\*). To show the approximation result, we distinguish three cases depending on the size of job  $j$ .

CASE 1. If  $p_j \leq \frac{1}{2}\text{OPT}'$ , Option 0 yields a schedule of makespan at most  $\frac{3}{2}\text{OPT}'$  by Corollary 2.1.

CASE 2. If  $p_j \geq \frac{3}{4}\text{OPT}'$ , then  $\frac{4}{3}p_j \geq \text{OPT}$  so that for any machine  $i$  it is feasible to migrate all but the largest job due to property (\*). Since there are at most  $m - 1$  jobs in  $N$  of size larger than  $\frac{1}{2}\text{OPT}'$ , there is a machine  $i$  that only contains jobs of size at most  $\frac{1}{2}\text{OPT}'$ . Then Option  $i$  yields a schedule with makespan at most  $\frac{3}{2}\text{OPT}'$ . (Apply the same argument as in Case 1 for the reassignment of removed jobs.)

CASE 3. It remains to consider the situation when  $p_j = \frac{1}{2}\text{OPT}' + \delta$  for some  $0 < \delta < \frac{1}{4}\text{OPT}'$ . In the following a job in  $N$  is called *huge* if its size is strictly larger than  $\text{OPT}' - \delta$  and *tiny* if its size is strictly smaller than  $\delta$ . The subsets of huge and tiny jobs are denoted by  $N_{\text{H}}$  and  $N_{\text{T}}$ , respectively. Notice that in an optimal schedule for  $N$  a huge job can only share a machine with tiny jobs.

CLAIM. In the given schedule  $S$  for  $N$  there exists a machine  $i$  that does not contain huge jobs and  $p(S^{-1}(i) \setminus N_{\text{T}}) \leq \text{OPT}$ .

In an optimal schedule for  $N$  the jobs in  $N \setminus (N_{\text{H}} \cup N_{\text{T}})$  are scheduled on the  $m - |N_{\text{H}}|$  machines that do not contain a huge job. This yields  $p(N \setminus (N_{\text{H}} \cup N_{\text{T}})) \leq (m - |N_{\text{H}}|)\text{OPT}$ . By a simple averaging argument, among the  $m - |N_{\text{H}}|$  machines that do not contain a huge job in the given schedule  $S$  there is one as claimed above.

It remains to show that Option  $i$  yields a  $\frac{3}{2}$ -approximate schedule for  $N \cup \{j\}$ . Since  $p(S^{-1}(i) \setminus N_{\text{T}}) \leq \text{OPT}$ , there is at most one job of size strictly larger than  $\frac{1}{2}\text{OPT}$  on machine  $i$ . Consequently, every job removed during Option  $i$  has size at most  $\frac{1}{2}\text{OPT}$  and the reassignment of these jobs does not cause a problem with respect to the bound  $\frac{3}{2}\text{OPT}'$ . We still need to show that the load of machine  $i$  after assigning job  $j$  does not exceed  $\frac{3}{2}\text{OPT}'$ .

If there is a tiny job that is not removed from machine  $i$ , then the total size of removed jobs is larger than  $\frac{4}{3}p_j - \delta \geq p_j$  and we are done since the load of machine  $i$  in the given schedule  $S$  is at most  $\frac{3}{2}\text{OPT}$ . We can assume in the following that all tiny jobs are removed from machine  $i$ .

If  $S^{-1}(i) \setminus N_{\text{T}}$  contains only one job, we are done since this job is not huge such that together with job  $j$  it does not exceed the bound  $\frac{3}{2}\text{OPT}'$  on the load of machine  $i$ . Finally, if  $S^{-1}(i) \setminus N_{\text{T}}$  contains more than one job, then, as already argued above, the second largest job has size at most  $\frac{1}{2}\text{OPT} \leq \frac{4}{3}p_j$  and is therefore removed (remember that the jobs are considered in order of non-increasing size during the removal phase). Since this job is not tiny, it has size at least  $\delta$  and the total size of removed jobs is at least  $\delta + p(S^{-1}(i) \cap N_{\text{T}})$ . Therefore the load of machine  $i$  after the removal phase is at most  $p(S^{-1}(i) \setminus N_{\text{T}}) - \delta \leq \text{OPT} - \delta$  and we can afford to assign job  $j$  to machine  $i$ .  $\square$

In Section 4 we present a more sophisticated algorithm with migration factor 4 and competitive ratio  $\frac{4}{3}$ . Moreover, in Section 5 we present a specialized algorithm for the case of two machines with migration factor 1 and competitive ratio  $\frac{7}{6}$ ; and we also show that this result is tight, i.e., migration factor 1 does not allow for a smaller competitive ratio.

**Robustness** The scheduling strategies for minimizing the makespan discussed in this section (and in Section 4) are robust in the following sense. The only invariant that we require in their analyses is that before the arrival of a new job the current schedule is  $\alpha$ -approximate. Job  $j$  can then be incorporated yielding again an  $\alpha$ -approximate schedule. In other words, we do not require that the current schedule is carefully constructed so far, to maintain the competitiveness in the next round. Only for Procedure 2 the schedule must additionally satisfy that, on any machine, the load excluding the largest job is at most the optimum makespan. This is a rather mild assumption as even the simple list scheduling algorithm and almost all natural scheduling strategies ensure this.

**Negative Results** Theorems 3.1 and 3.2 raise the question which migration factor is really necessary to achieve competitive ratio  $\frac{3}{2}$ . We can prove that Procedure 2 is optimal in the sense that any robust scheduling strategy needs migration factor at least  $\frac{4}{3}$  in order to maintain competitive ratio  $\frac{3}{2}$ .

LEMMA 3.2 *There exists a  $\frac{3}{2}$ -approximate schedule fulfilling property (\*) such that, for any  $0 < \epsilon < \frac{4}{21}$ ,*




Figure 1: An instance with 26 machines and 50 jobs. The load of the machines is depicted on the vertical axis. On the right hand side an optimal schedule with makespan 8 is given. On the left hand side there is a  $\frac{3}{2}$ -approximate schedule fulfilling property (\*). If a new job of size  $6 + \epsilon$  arrives, jobs of total size at least 8 have to be moved in order to obtain a schedule that is still  $\frac{3}{2}$ -approximate.

upon arrival of a new job, migration factor  $\frac{4}{3} - \epsilon$  is necessary to achieve  $\frac{3}{2}$ -competitiveness.

**PROOF.** The situation is depicted in Figure 1. There are 26 machines and 50 jobs. A  $\frac{3}{2}$ -approximate schedule fulfilling property (\*) with makespan 12 is given on the left hand side of the figure while an optimal schedule with makespan 8 is given on the right hand side. In both schedules, machines 3 to 10 are identically packed and the same holds for machines 11 to 26. Notice that machine 2 is empty in the optimal schedule. Therefore, upon arrival of a new job of size  $\frac{24}{4-3\epsilon}$ , the optimal makespan remains 8. Notice that  $6 < \frac{24}{4-3\epsilon} < 7$  due to our choice of  $\epsilon$ .

In order to incorporate the new job into the schedule on the left hand side of Figure 1 without increasing the makespan, jobs of total size at least 8 have to be migrated: First of all notice that it does not make sense to assign the new job to one of the machines currently containing a job of size 7. If the new job is assigned to one of the machines currently containing a job of size 6, this job has to be removed from the machine. In order to reassign it without increasing the makespan, two more jobs of size 1 have to be moved. Finally, in order to schedule the new job on one of the first two machines, jobs of total size at least 8 have to be moved.

Thus, the migration factor is at least  $8 \cdot \frac{4-3\epsilon}{24} = \frac{4}{3} - \epsilon$ .  $\square$

An additional feature of Procedures 1 and 2 is that they are *local* in the sense that they migrate jobs only from the machine where the newly arrived job is assigned to.

**LEMMA 3.3** *There is a class of optimal schedules for which, upon arrival of a new job, it is not possible to achieve a competitive ratio smaller than  $\frac{3}{2}$  using only local migration (even if an arbitrary migration factor is allowed).*

**PROOF.** The following optimal schedule on  $m$  machines, upon the arrival of a new job, enforces a competitive ratio of at least  $3/(2 + \frac{2}{m})$  for any amount of local migration. This bound converges to  $\frac{3}{2}$  for large  $m$ . Machines 1 and 2 each contain one job of size  $\frac{1}{2}$  and  $\frac{m}{2}$  jobs of size  $\frac{1}{m}$  (we assume that  $m$  is even). All other machines contain a single job of size 1; see Figure 2. The newly arriving job has size 1. The optimum makespan is  $1 + \frac{1}{m}$  and the makespan achievable by any local strategy is  $\frac{3}{2}$  (by scheduling the new job e.g. on machine 1 and migrating all small jobs to other machines).  $\square$

**4. A  $\frac{4}{3}$ -Competitive Strategy with Migration Factor 4** Taking up the line of research pursued in Section 3, we show that competitive ratio  $\frac{4}{3}$  can be achieved by a more sophisticated algorithm with migration factor 4.

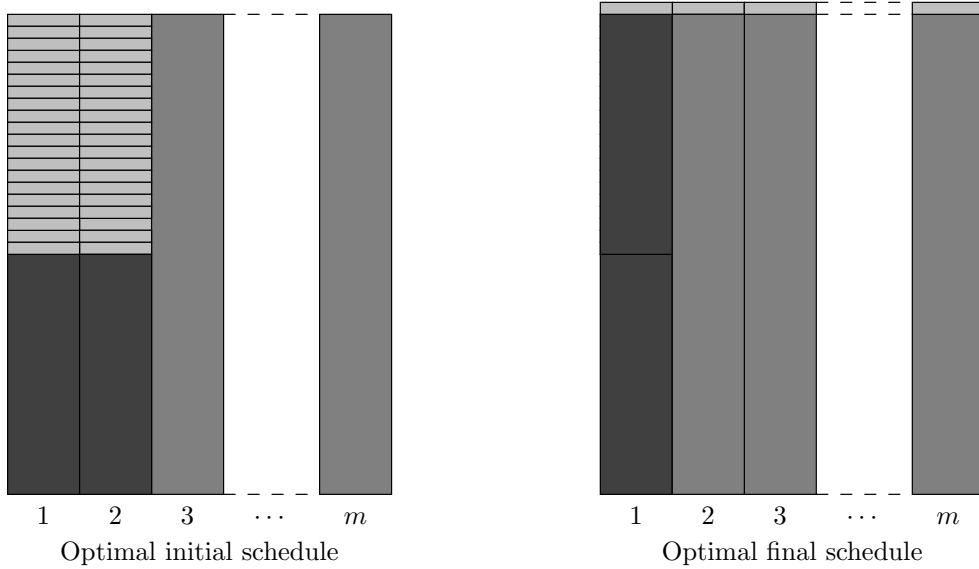


Figure 2: An instance where local migration is not better than  $\frac{3}{2}$ -competitive.

PROCEDURE 3

Input: A schedule  $S$  for a set of jobs  $N$  and a new job  $j \notin N$ .

Output: A schedule for  $N \cup \{j\}$ .

Choose one of the following  $m + 1$  options which minimizes the resulting makespan.

Option 0: Assign job  $j$  to a least loaded machine.

Option  $i$  [for  $i \in \{1, \dots, m\}$ ]: If  $p_j < \max\{\frac{1}{2}p(S^{-1}(i)), p_{\max}(S^{-1}(i))\}$ , then skip Phase 1 and set  $\ell := j$ .

PHASE 1: Remove all jobs from machine  $i$  and schedule job  $j$  there. Assign the removed jobs successively to a least loaded machine, except for the largest job  $\ell$  which is handled in Phase 2.

PHASE 2: Choose one of the following  $m + 1$  sub-options which minimizes the resulting makespan.

Sub-option 0: Assign job  $\ell$  to a least loaded machine.

Sub-option  $k$  [for  $k \in \{1, \dots, m\}$ ]: Ignoring the largest job on machine  $k$ , repeatedly remove jobs from this machine; stop before the total size of removed jobs exceeds  $2p_\ell$ . Assign job  $\ell$  to machine  $k$ . Assign the removed jobs successively to a least loaded machine.

THEOREM 4.1 Given a  $\frac{4}{3}$ -approximate schedule for  $N$  and a new job  $j \notin N$ , Procedure 3 outputs a  $\frac{4}{3}$ -approximate schedule for  $N \cup \{j\}$ . The migration factor is bounded by 4.

PROOF. We first prove that the migration factor in Option  $i > 0$  is bounded by 4. If Phase 1 is not skipped, then the total processing time of jobs migrated in Phase 1 is  $p(S^{-1}(i)) \leq 2p_j$ . Moreover, the total processing time of jobs migrated in Phase 2 is bounded by  $2p_\ell$ . Notice that  $p_\ell = p_j$  or  $p_\ell = p_{\max}(S^{-1}(i)) \leq p_j$ . Thus, the total processing time of migrated jobs is bounded by  $4p_j$ .

We partition the set of jobs  $N \cup \{j\}$  into subsets of *small*, *medium* and *large* jobs. A job  $j'$  is *small* if  $p_{j'} \leq \frac{1}{3}\text{OPT}'$  and it is *large* if  $p_{j'} > \frac{2}{3}\text{OPT}'$ . Otherwise, if  $\frac{1}{3}\text{OPT}' < p_{j'} \leq \frac{2}{3}\text{OPT}'$  we say that job  $j'$  is *medium*. If job  $j$  is small, then Option 0 yields a  $\frac{4}{3}$ -approximate schedule by Corollary 2.1. We can therefore assume from now on that  $j$  is medium or large.

In an optimal schedule for  $N \cup \{j\}$  a large job can share a machine only with small jobs. Since  $j$  is medium or large, the number  $m'$  of large jobs in  $N$  is bounded by  $m - 1$ . Moreover, in an optimal schedule for  $N \cup \{j\}$  at most two medium jobs can share a machine such that the number of medium jobs in  $N$  is at most  $2(m - m') - 1$  (the  $-1$  is due to the fact that  $j$  is medium or large). Thus there exists a machine  $i$  that contains at most one medium job and no large jobs in the given schedule for  $N$ . We show that Option  $i$  yields a  $\frac{4}{3}$ -approximate schedule for  $N \cup \{j\}$ .



We first argue that the schedule for  $(N \cup \{j\}) \setminus \{\ell\}$  after Phase 1 has makespan at most  $\frac{4}{3}\text{OPT}'$ . If Phase 1 is being skipped, this is obvious. Otherwise, job  $j$  is assigned to the empty machine  $i$ . By choice of  $i$  the jobs in  $S^{-1}(i) \setminus \{\ell\}$  are small such that assigning them successively to a least loaded machine does not cause a problem with respect to the bound  $\frac{4}{3}\text{OPT}'$  by Corollary 2.1.

It remains to consider Phase 2 where a schedule  $S'$  for  $(N \cup \{j\}) \setminus \{\ell\}$  with makespan at most  $\frac{4}{3}\text{OPT}'$  and job  $\ell$  are given. We first argue that job  $\ell$  is medium or small. If Phase 1 has been skipped, then  $\ell = j$  and  $p_\ell = p_j < \frac{1}{2}p(S^{-1}(i)) \leq \frac{2}{3}\text{OPT}'$ . Otherwise, if Phase 1 has not been skipped, job  $\ell$  is medium or small by choice of  $i$ .

If  $\ell$  is small, Sub-option 0 yields a  $\frac{4}{3}$ -approximate schedule by Corollary 2.1. We can therefore assume from now on that  $\ell$  is medium.

By the same argument already used above for schedule  $S$ , there exists a machine  $i'$  such that  $S'^{-1}(i')$  contains at most one medium job and no large jobs. We show that Sub-option  $i'$  yields a  $\frac{4}{3}$ -approximate schedule for  $N \cup \{j\}$ .

Let  $\ell'$  denote the largest job on machine  $i'$ . Since all jobs removed from machine  $i'$  in Sub-option  $i'$  are small, assigning them successively to a least loaded machine does not cause a problem with respect to the bound  $\frac{4}{3}\text{OPT}'$ . It therefore remains to consider the situation on machine  $i'$  after assigning job  $\ell$  to it. If all jobs except  $\ell'$  are removed from  $i'$ , then assigning  $\ell$  yields a total load of  $p_{\ell'} + p_\ell \leq \frac{4}{3}\text{OPT}'$  (since  $\ell$  is medium and  $\ell'$  is medium or small). Otherwise, if apart from  $\ell'$  other (small) jobs remain on machine  $i'$ , then the total size of jobs removed is larger than  $2p_{\ell'} - \frac{1}{3}\text{OPT}' > p_\ell$  (since  $\ell$  is medium). Therefore the total load on machine  $i'$  after assigning job  $\ell$  is at most  $\frac{4}{3}\text{OPT}'$ .  $\square$

**5. The Two Machine Case** In this section we present an online algorithm for the two machine case which achieves competitive ratio  $\frac{7}{6}$  with migration factor 1. Before we discuss the algorithm, we first prove that this result is best possible.

**THEOREM 5.1** *Any online algorithm with migration factor at most 1 has competitive ratio at least  $\frac{7}{6}$ .*

**PROOF.** Consider an instance consisting of four jobs with sizes  $\frac{1}{3}$ ,  $\frac{1}{3}$ ,  $\frac{1}{2}$ , and  $\frac{1}{2}$ . The optimal makespan is obviously  $\text{OPT} = \frac{5}{6}$ . It is easy to verify that the only  $\frac{7}{6}$ -approximate schedule is the optimal schedule that assigns one job of size  $\frac{1}{3}$  and one job of size  $\frac{1}{2}$  to each machine. After the arrival of a new job  $j$  of size  $p_j = \frac{1}{3}$ , the optimal makespan is  $\text{OPT}' = 1$ . With migration factor 1, however, one can only achieve makespan  $\frac{7}{6}$ .  $\square$

Before we present our particular online algorithm, we first prove the *existence* of an online algorithm with competitive ratio  $\frac{7}{6}$  and migration factor 1. We even show that there is such an online algorithm which is robust and uses only local migration; see Section 3 for the definitions of robustness and local migration.

**LEMMA 5.1** *Consider an arbitrary  $\frac{7}{6}$ -approximate schedule for a given set of jobs  $N$  on two machines and a newly arriving job  $j$ . There is a subset of jobs  $X$  of total size  $p(X) \leq p_j$  residing on one of the two machines such that scheduling job  $j$  on this machine and migrating the jobs in  $X$  to the other machine yields a  $\frac{7}{6}$ -approximate schedule for  $N \cup \{j\}$ .*

**PROOF.** If  $p_j \leq \frac{1}{3}\text{OPT}'$ , it follows from Corollary 2.1 that assigning job  $j$  to a least loaded machine yields the desired schedule, i.e.,  $X$  can be chosen to be the empty set in this case. In the following we assume that  $p_j > \frac{1}{3}\text{OPT}'$ .

Consider an optimal schedule for  $N \cup \{j\}$ . We describe the difference between this optimal schedule and the given schedule for  $N$ ; see Figure 3. Let  $\delta_1$  be the subset of jobs assigned to the first machine in both schedules. Similarly,  $\Delta_2$  is the set of jobs assigned to the second machine in both schedules. The remaining two sets  $\delta_2$  and  $\Delta_1$  capture the differences between these two schedules except for job  $j$ , which is only present in the optimal schedule; see Figure 3. In the following we assume without loss of generality that  $p(\delta_1) \leq p(\delta_2)$ . Since

$$p(\delta_1) + p(\delta_2) + p_j \leq \text{OPT}' , \quad (2)$$

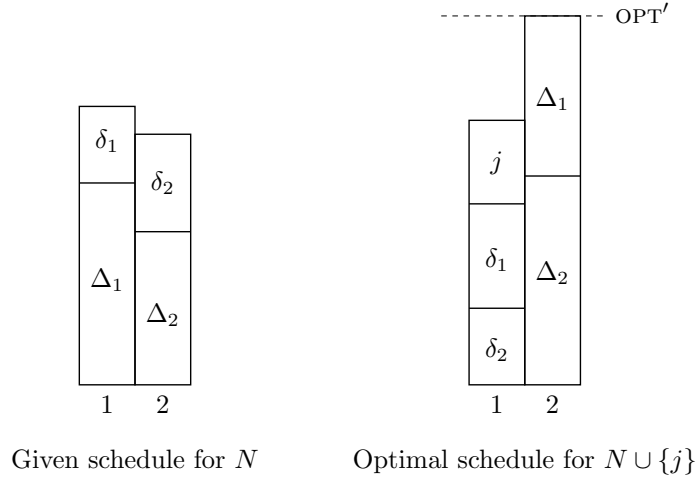


Figure 3: Comparison of given schedule for  $N$  and optimal schedule for  $N \cup \{j\}$  on two machines.

we get  $p(\delta_1) < \frac{1}{3}\text{OPT}'$  as  $p_j > \frac{1}{3}\text{OPT}'$ . Also notice that

$$p(\Delta_1) + p(\Delta_2) \leq \text{OPT}' . \quad (3)$$

Without exceeding migration factor 1, we can remove the subset of jobs  $\delta_1$  from the first machine, assign  $j$  to this machine, and then assign all jobs in  $\delta_1$  to the least loaded machine. If  $p_j + p(\Delta_1) \leq \frac{7}{6}\text{OPT}'$ , we are done by Corollary 2.1; set  $X := \delta_1$  or  $X := \emptyset$ , depending on which machine has smaller load after  $j$  has been scheduled. We thus assume from now on that

$$p_j + p(\Delta_1) > \frac{7}{6}\text{OPT}' \quad (4)$$

and thus, by (3),

$$p_j > p(\Delta_2) + \frac{1}{6}\text{OPT}' . \quad (5)$$

We distinguish two cases.

CASE 1:  $p(\Delta_2) \leq \frac{1}{3}\text{OPT}'$ . Since  $p_j + \delta_2 \leq \text{OPT}'$  by (2), removing the jobs in  $\Delta_2$  from the second machine, assigning job  $j$  to this machine, and then assigning all jobs in  $\Delta_2$  to the least loaded machine yields the desired result by Corollary 2.1; set  $X := \Delta_2$  or  $X := \emptyset$ , depending on which machine has smaller load after  $j$  has been scheduled.

CASE 2:  $p(\Delta_2) > \frac{1}{3}\text{OPT}'$ . Then  $p_j > \frac{1}{2}\text{OPT}'$  by (5) and thus  $p(\delta_1) + p(\delta_2) < \frac{1}{2}\text{OPT}'$  by (2). Since  $p(\Delta_1) < \frac{2}{3}\text{OPT}'$  by (3), we get

$$p(\Delta_1) + p(\delta_1) + p(\delta_2) < \frac{7}{6}\text{OPT}' . \quad (6)$$

If  $p_j + p(\Delta_2) \leq \frac{7}{6}\text{OPT}'$ , then  $X := \delta_2$  is a feasible choice and we are done. Otherwise,  $p_j + p(\Delta_2) > \frac{7}{6}\text{OPT}'$  together with (4) and (3) implies that  $p_j > \frac{2}{3}\text{OPT}'$ . As a consequence of (2) we get  $p(\delta_1) + p(\delta_2) < \frac{1}{3}\text{OPT}'$ . Moreover, since  $p(\delta_1) \leq p(\delta_2)$  we obtain  $p(\delta_1) < \frac{1}{6}\text{OPT}'$ . By (3) we get  $p(\Delta_1) + p(\Delta_2) + p(\delta_1) < \frac{7}{6}\text{OPT}'$ . Therefore  $X := \Delta_2$  is a feasible choice and the proof is complete.  $\square$

Of course, the result in Lemma 5.1 can be easily turned into an online algorithm with competitive ratio  $\frac{7}{6}$  and migration factor 1 using brute force: Simply determine the subset of jobs  $X$  by complete enumeration. The following more efficient algorithm is based on the intuition that it is sufficient to know those jobs in  $X$  whose size is at least  $\frac{1}{3}\text{OPT}'$ . The remaining jobs can be chosen greedily. We call a job in the initial schedule  $S$  *large* if it is among the three largest jobs on its machine. Otherwise, it is *small*. Since in a  $\frac{7}{6}$ -approximate schedule for  $N$  there are at most three jobs of size at least  $\frac{1}{3}\text{OPT}'$  on each machine, we know that small jobs have size at most  $\frac{1}{3}\text{OPT}'$ . For  $i \in \{1, 2\}$ , we define a family of subsets of  $N$  by

$$\mathcal{L}_i := \{Y \subseteq S^{-1}(i) \mid p(Y) \leq p_j \text{ and } Y \text{ only contains large jobs}\} .$$

As there are at most three large jobs on machine  $i$  in schedule  $S$ , the family  $\mathcal{L}_i$  contains at most  $2^3$  subsets of jobs.

## PROCEDURE 4

Input: A schedule  $S$  for a set of jobs  $N$  and a new job  $j \notin N$ .

Output: A schedule for  $N \cup \{j\}$ .

We define an option for each  $i = 1, 2$  and  $Y \in \mathcal{L}_i$ . Choose one of these (at most 16) options which minimizes the resulting makespan.

Option  $i, Y$ : Migrate all jobs in  $Y$  to the other machine and assign job  $j$  to machine  $i$ . Consider the small jobs in  $S^{-1}(i)$  in arbitrary order: Migrate a small job unless the total size of migrated jobs will exceed  $p_j$  or migrating the job will not lead to an improved makespan.

**THEOREM 5.2** Given a  $\frac{7}{6}$ -approximate schedule for  $N$  and a new job  $j \notin N$ , Procedure 4 outputs a  $\frac{7}{6}$ -approximate schedule for  $N \cup \{j\}$ . The migration factor is at most 1.

**PROOF.** The bound on the migration factor is clear from the description of Procedure 4. Let  $i \in \{1, 2\}$  and  $X \subseteq S^{-1}(i)$  be a subset as in Lemma 5.1. Moreover, let  $Y \subseteq X$  be the subset of large jobs in  $X$ . We prove that Option  $i, Y$  yields a  $\frac{7}{6}$ -approximate schedule for  $N \cup \{j\}$ . We distinguish two cases.

**CASE 1:**  $X = \emptyset$ . In this case, assigning job  $j$  to machine  $i$  yields a  $\frac{7}{6}$ -approximate schedule for  $N \cup \{j\}$ . In Option  $i, \emptyset$ , our algorithm assigns  $j$  to machine  $i$  and afterwards only migrates small jobs if this decreases the makespan. Therefore it also outputs a  $\frac{7}{6}$ -approximate schedule for  $N \cup \{j\}$ .

**CASE 2:**  $X \neq \emptyset$ . In this case we may assume that assigning job  $j$  to any of the two machines yields a schedule with makespan strictly larger than  $\frac{7}{6}\text{OPT}'$  (otherwise we can choose  $X := \emptyset$  and are back in Case 1). In other words,

$$p(S^{-1}(1)) > \frac{7}{6}\text{OPT}' - p_j \quad \text{and} \quad p(S^{-1}(2)) > \frac{7}{6}\text{OPT}' - p_j. \quad (7)$$

We refer to machine  $i$  in the following as the *first machine* and to the other machine as the *second machine*. We first argue that the load of the second machine in the schedule computed by Option  $i, Y$  is at most  $\frac{7}{6}\text{OPT}'$ . Since  $Y \subseteq X$ , migrating the large jobs in  $Y$  cannot increase the load of the second machine beyond  $\frac{7}{6}\text{OPT}'$ . Thus, before we start to migrate small jobs, the load of the second machine is at most  $\frac{7}{6}\text{OPT}'$ . It follows from Corollary 2.1 that the load of the second machine remains below  $\frac{7}{6}\text{OPT}'$  if we improve the makespan by migrating small jobs of size at most  $\frac{1}{3}\text{OPT}'$ .

It remains to show that the load of the first machine  $i$  in the schedule computed by Option  $i, Y$  is at most  $\frac{7}{6}\text{OPT}'$ . We distinguish two sub-cases.

**CASE 2a:** All small jobs in  $S^{-1}(i)$  are migrated to the second machine in Option  $i, Y$ . In this case we have migrated all jobs in  $X$  and possibly some more. In particular, the load of machine  $i$  in our schedule is at most the load of machine  $i$  in the  $\frac{7}{6}$ -approximate schedule described in Lemma 5.1.

**CASE 2b:** There is a small job  $k \in S^{-1}(i)$  which is not migrated in Option  $i, Y$ . There can be two reasons for not migrating job  $k$ . If  $k$  is not migrated because this does not decrease the makespan, then the load of machine  $i$  is at most  $\frac{7}{6}\text{OPT}'$  by Corollary 2.1 and we are done. Otherwise, the total size of jobs that have already been migrated is larger than  $p_j - p_k \geq p_j - \frac{1}{3}\text{OPT}'$ . Because of (7) the load of the second machine is at least  $\frac{7}{6}\text{OPT}' - p_j + p_j - \frac{1}{3}\text{OPT}' = \frac{5}{6}\text{OPT}'$ . As the total load of the two machines is at most  $2\text{OPT}'$ , the load of machine  $i$  is at most  $\frac{7}{6}\text{OPT}'$ .  $\square$

**6. An Online Approximation Scheme with Constant Migration** The results presented in Section 3, 4, and 5 raise the question how far the competitive ratio for online algorithms with constant migration factor can be decreased. We first prove that optimality (i.e., competitive ratio 1) cannot be achieved. However, for any fixed  $\epsilon > 0$  we can get down to competitive ratio  $1 + \epsilon$ .

The following lemma states that online algorithms with constant migration factor and competitive ratio 1 do not exist.

**LEMMA 6.1** Any online algorithm which always maintains an optimal solution has migration factor  $\Omega(m)$ .

**PROOF.** Consider a scheduling instance with  $m$  machines and  $2m - 2$  jobs, two of size  $\frac{i}{m}$  for all  $i = 1, \dots, m - 1$ . Up to permutations of machines, any optimum schedule has the structure depicted




Figure 4: An instance where all machine configurations have to change to maintain optimality.

in the left part of Figure 4. The optimum makespan is  $1 - \frac{1}{m}$ . When a new job of size 1 arrives, the optimum makespan increases to 1. Again, the structure of an optimum schedule for the enlarged instance is unique; see the right hand side of Figure 4. From each machine in  $\{2, \dots, m-1\}$ , at least one of the two jobs has to move. Hence the total size of jobs that have to be migrated is at least

$$\frac{1}{m} \sum_{i=1}^{m-2} \min\{i, m-1-i\} \geq \frac{1}{m} \sum_{i=1}^{\lfloor \frac{m-2}{2} \rfloor} i$$

which is in  $\Omega(m)$ .  $\square$

In the remainder of this section,  $\epsilon > 0$  is a fixed constant. (In the following we assume without loss of generality that  $\epsilon < 1$ .) We develop an algorithm with competitive ratio  $1 + O(\epsilon)$  and constant migration factor  $\beta(\epsilon)$ . This algorithm is based on a combination of several techniques known from the area of polynomial time approximation schemes for machine scheduling problems. The root of the presented method is the polynomial time approximation scheme by Hochbaum and Shmoys [18] and its refinements [17]. The pivotal point of our online approach is to establish a connection to classical sensitivity analysis for integer linear programming.

The following observation belongs to the folklore in the field of scheduling; see, e.g., [1].

**OBSERVATION 6.1** *Rounding up each job's processing time to the nearest integer power of  $1 + \epsilon$  increases the makespan of an arbitrary schedule at most by a factor of  $1 + \epsilon$ . In particular, in specifying a  $(1 + O(\epsilon))$ -competitive algorithm we can assume that all processing times are integer powers of  $1 + \epsilon$ .*

In the remainder of this section we thus consider the situation when all processing times are integer powers of  $1 + \epsilon$ . The current set of jobs is denoted by  $N$  and the newly arriving job by  $j$ . Moreover we set  $N' := N \cup \{j\}$ . A job in  $N'$  is called *large* if its processing time is at least  $\epsilon \text{LB}(N')$  where, according to Observation 2.1,

$$\text{LB}(N') := \max \left\{ \frac{p(N')}{m}, p_{\max}(N') \right\} ;$$

otherwise, it is called *small*. The subsets of large and small jobs in  $N'$  are denoted by  $N'_L$  and  $N'_S$ , respectively. Notice that  $N'$  is the disjoint union of  $N'_L$  and  $N'_S$ . We also define a much finer partition of the jobs in  $N'$  into classes of equal-size jobs  $N'_i$ ,  $i \in \mathbb{Z}$ , with

$$N'_i := \{j' \in N' \mid p_{j'} = (1 + \epsilon)^i\} .$$

Moreover, let  $N_L := N'_L \cap N$ ,  $N_S := N'_S \cap N$ , and  $N_i := N'_i \cap N$  for  $i \in \mathbb{Z}$ . Then  $N$  is the disjoint union of  $N_L$  and  $N_S$ ; also the subsets  $N_i$ ,  $i \in \mathbb{Z}$ , form a partition of  $N$ .

It is a well known fact from the area of approximation algorithms for machine scheduling problems that the crucial skeleton of a schedule is given by the assignment of large jobs. Due to Corollary 2.1 the

small jobs can later be filled in greedily. We start by describing the set of all schedules of large jobs as the set of integral lattice points of a polyhedron.

Let  $I := \{i \in \mathbb{Z} \mid \epsilon \text{LB}(N') \leq (1 + \epsilon)^i \leq p_{\max}(N')\}$  denote the index set of those classes that might contain large jobs such that  $N'_L = \bigcup_{i \in I} N'_i$  and  $N_L = \bigcup_{i \in I} N_i$ . We can bound the number of such job classes as follows.

**OBSERVATION 6.2** *The number  $|I|$  of classes of large jobs is at most  $O(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$ .*

**PROOF.** By definition of  $I$  it holds that

$$|I| \leq 1 + \log_{1+\epsilon} \frac{p_{\max}(N')}{\epsilon \text{LB}(N')} \leq 1 + \log_{1+\epsilon} \frac{1}{\epsilon} \leq 1 + \frac{2}{\epsilon} \log \frac{1}{\epsilon} \in O\left(\frac{1}{\epsilon} \log \frac{1}{\epsilon}\right) .$$

This concludes the proof.  $\square$

Given an assignment of large jobs to machines, we say that a particular machine obeys *configuration*  $k : I \rightarrow \mathbb{N}_0$  if, for all  $i \in I$ , exactly  $k(i)$  jobs of size  $(1 + \epsilon)^i$  are assigned to this machine. The set of configurations that can occur in any schedule for  $N'_L$  is denoted by

$$\mathbf{K} := \{k : I \rightarrow \mathbb{N}_0 \mid k(i) \leq |N'_i| \text{ for all } i \in I\} .$$

Up to permutations of machines and up to permutations of equal size jobs, an arbitrary schedule for  $N'_L$  can be described by specifying, for each  $k \in \mathbf{K}$ , the number  $y_k$  of machines that obey configuration  $k$ . Conversely, a vector  $y \in \mathbb{N}_0^{\mathbf{K}}$  specifies a feasible  $m$ -machine-schedule for  $N'_L$  if and only if

$$\sum_{k \in \mathbf{K}} y_k = m , \tag{8}$$

i.e., the schedule uses exactly  $m$  machines, and

$$\sum_{k \in \mathbf{K}} k(i) y_k = |N'_i| \quad \text{for all } i \in I, \tag{9}$$

i.e., the set of jobs assigned to the machines is  $N'$ . We denote the set of vectors  $y \in \mathbb{N}_0^{\mathbf{K}}$  satisfying (8) and (9) by  $\mathbf{S}'$ . Thus,  $\mathbf{S}'$  represents the set of all schedules for  $N'_L$  (up to permutations of machines and up to permutations of equal size jobs). If we replace  $|N'_i|$  with  $|N_i|$  on the right hand side of (9), then constraints (8) and (9) characterize the set  $\mathbf{S}$  of all schedules for  $N_L$ .

For a configuration  $k \in \mathbf{K}$  let

$$\text{load}(k) := \sum_{i \in I} (1 + \epsilon)^i k(i)$$

denote the load of a machine obeying configuration  $k$ . The makespan of a schedule  $y \in \mathbf{S}$  or  $y \in \mathbf{S}'$  is thus equal to  $\max\{\text{load}(k) \mid y_k > 0\}$ . For  $\mu \geq 0$  we denote the set of configurations of load at most  $\mu$  by

$$\mathbf{K}(\mu) := \{k \in \mathbf{K} \mid \text{load}(k) \leq \mu\} .$$

The subsets of all schedules for  $N_L$  and  $N'_L$  with makespan at most  $\mu$  are denoted by

$$\mathbf{S}(\mu) := \{y \in \mathbf{S} \mid y_k = 0 \text{ if } \text{load}(k) > \mu\}$$

and

$$\mathbf{S}'(\mu) := \{y \in \mathbf{S}' \mid y_k = 0 \text{ if } \text{load}(k) > \mu\} ,$$

respectively.

In the following, we usually interpret a schedule  $y \in \mathbf{S}(\mu)$  or  $y \in \mathbf{S}'(\mu)$  as a vector in  $\mathbb{N}_0^{\mathbf{K}(\mu)}$  by ignoring all zero-entries corresponding to configurations in  $\mathbf{K} \setminus \mathbf{K}(\mu)$ . We can thus write

$$\mathbf{S}(\mu) = \{y \in \mathbb{N}_0^{\mathbf{K}(\mu)} \mid A(\mu) y = b\}$$

and

$$\mathbf{S}'(\mu) = \{y \in \mathbb{N}_0^{\mathbf{K}(\mu)} \mid A(\mu) y = b'\} ,$$

where  $A(\mu)$  is a matrix in  $\mathbb{N}_0^{(1+|I|) \times |\mathbf{K}(\mu)|}$  and  $b, b'$  are vectors in  $\mathbb{N}_0^{1+|I|}$ . The first row of the linear system  $A(\mu) y = b'$  corresponds to constraint (8); the remaining  $|I|$  rows correspond to constraints (9).

The next observation is crucial for applying sensitivity analysis of integer linear programming later.

OBSERVATION 6.3 *The vectors  $b$  and  $b'$  are equal for all but at most one entry where they differ by 1.*

PROOF. The first entry of both  $b$  and  $b'$  is equal to  $m$  (see the right hand side of (8)). The remaining entries are equal to  $|N_i|$  and  $|N'_i|$ , respectively. Notice that  $|N_i| = |N'_i|$  unless  $j \in N'_i$ . In the latter case  $|N'_i| = |N_i| + 1$ .  $\square$

The next lemma contains the key insight for the result presented in this section.

LEMMA 6.2 *Consider a schedule for  $N_L$  and denote its makespan by  $\mu$ . This schedule can be turned into a schedule for  $N'_L$  with makespan at most  $\max\{\mu, \text{OPT}(N'_L)\}$  while touching at most  $2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}$  machines.*

PROOF. Let  $\mu' := \max\{\mu, \text{OPT}(N'_L)\}$ . If  $N_L = N'_L$  (i.e., if  $j$  is small), there is nothing to be shown. Moreover, if  $\mu' \geq 2 \text{OPT}(N'_L)$ , the required schedule can be obtained by simply assigning  $j$  to the least loaded machine. Notice that the load of this machine as well as the size of job  $j$  are at most  $\text{OPT}(N'_L)$ . We can therefore assume in the following that  $j \in N'_L$  and  $\mu < 2 \text{OPT}(N'_L)$ .

Let  $y \in \mathbf{S}(\mu')$  denote the given schedule for  $N_L$ . Then,  $y$  satisfies

$$A(\mu')y = b, \quad y \in \mathbb{N}_0^{\mathbf{K}(\mu')}. \quad (10)$$

We are looking for a schedule  $y' \in \mathbf{S}'(\mu')$ , that is,  $y'$  must satisfy

$$A(\mu')y' = b', \quad y' \in \mathbb{N}_0^{\mathbf{K}(\mu')}. \quad (11)$$

By Observation 6.3, the right hand sides of the systems of linear equations in (10) and (11) are equal for all but one entry where they differ by 1. Using a sensitivity result from integer linear programming<sup>2</sup> there exists a solution  $y'$  to (11) satisfying

$$\|y - y'\|_\infty \leq 3 |\mathbf{K}(\mu')| \Delta,$$

where  $\Delta$  is an upper bound on the absolute value of any sub-determinant of the matrix  $A(\mu')$ . The number of machines whose configurations are different in schedules  $y$  and  $y'$  is bounded from above by

$$\|y - y'\|_1 \leq |\mathbf{K}(\mu')| \cdot \|y - y'\|_\infty \leq 3 |\mathbf{K}(\mu')|^2 \Delta, \quad (12)$$

To complete the proof, we have to show that the right hand side of (12) is constant. First we give an upper bound on the number of configurations  $|\mathbf{K}(\mu')|$ , i.e., on the number of machine configurations with load at most  $\mu'$ . Since each job in  $N'_L$  has size at least

$$\epsilon \text{LB}(N') \stackrel{(1)}{\geq} \frac{\epsilon}{2} \text{OPT}(N') \geq \frac{\epsilon}{2} \text{OPT}(N'_L) > \frac{\epsilon}{4} \mu',$$

there are less than  $\frac{4}{\epsilon}$  jobs in any configuration  $k \in \mathbf{K}(\mu')$  with load at most  $\mu'$ ; in particular, it holds that  $k(i) < \frac{4}{\epsilon}$  for all  $i \in I$ . This yields

$$|\mathbf{K}(\mu')| \leq \left(\frac{4}{\epsilon}\right)^{|I|} = 2^{|I| \log \frac{4}{\epsilon}} \in 2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})} \quad (13)$$

by Observation 6.2.

We now turn to  $\Delta$  which must be an upper bound on the absolute value of any sub-determinant of the matrix  $A(\mu')$ . The entries in the first row of  $A(\mu')$  are all 1 (see (8)) and the remaining entries are of the form  $k(i) < \frac{4}{\epsilon}$  (see (9)). The maximum dimension of a square sub-matrix of  $A(\mu')$  is at most the number of rows which is  $1 + |I|$ . Hence the absolute value of any sub-determinant is upper-bounded by

$$\Delta := (1 + |I|)! \left(\frac{4}{\epsilon}\right)^{|I|} \leq ((1 + |I|)\frac{4}{\epsilon})^{|I|} = 2^{|I|(\log(1 + |I|) + \log \frac{4}{\epsilon})} \in 2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}. \quad (14)$$

The claimed result follows by plugging the upper bound (13) on  $|\mathbf{K}(\mu')|$  and the upper bound (14) on  $\Delta$  into (12).  $\square$

We now return to the complete set of jobs  $N'$  also including small jobs. Building upon Lemma 6.2 we show how to incorporate small jobs in the following theorem.

<sup>2</sup>For the convenience of the reader we quote the used result:

COROLLARY 17.2A (from [28]). Let  $A$  be an integral  $m \times n$ -matrix, such that each subdeterminant of  $A$  is at most  $\Delta$  in absolute value, let  $b'$  and  $b''$  be column  $m$ -vectors, and let  $c$  be a row  $n$ -vector. Suppose  $\max\{cx \mid Ax \leq b'; x \text{ integral}\}$  and  $\max\{cx \mid Ax \leq b''; x \text{ integral}\}$  are finite. Then for each optimum solution  $z'$  of the first maximum there exists an optimum solution  $z''$  of the second maximum such that  $\|z' - z''\|_\infty \leq n\Delta(\|b' - b''\|_\infty + 2)$ .

**THEOREM 6.1** *Any  $(1 + \epsilon)$ -approximate schedule for  $N$  can be turned into a  $(1 + \epsilon)$ -approximate schedule for  $N' = N \cup \{j\}$  such that the total size of jobs that have to be moved is bounded by a constant  $\beta(\epsilon)$  times  $p_j$  with  $\beta(\epsilon) \in 2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}$ .*

**PROOF.** We distinguish two cases. If the newly arrived job is small, i.e.,  $p_j < \epsilon \text{LB}(N')$ , then  $j$  can simply be assigned to a least loaded machine by Corollary 2.1 and no job in  $N$  has to be moved.

It remains to consider the situation when job  $j$  is large. The given schedule for  $N$  induces a schedule for  $N_L$  with makespan  $\mu \leq (1 + \epsilon) \text{OPT}(N) \leq (1 + \epsilon) \text{OPT}(N')$ . By Lemma 6.2, the latter schedule can be turned into a schedule for  $N'_L$  with makespan at most

$$\max\{\mu, \text{OPT}(N'_L)\} \leq (1 + \epsilon) \text{OPT}(N') ,$$

by touching only  $2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}$  machines. In the following, this subset of machines of constant size is denoted by  $M'$ . We construct a schedule for  $N'$  as follows:

- (i) Start with the schedule for  $N'_L$  discussed above.
- (ii) The small jobs in  $N_S$  which the given schedule for  $N$  assigned to one of the machines in  $M \setminus M'$  are assigned to the same machine again.
- (iii) The remaining jobs in  $N_S = N'_S$  are assigned one after another to a least loaded machine.

The makespan of the partial schedule constructed in steps (i) and (ii) is bounded by the maximum of the makespan of the given schedule for  $N$  and the optimal makespan of the schedule for  $N'_L$ . It is thus bounded by  $(1 + \epsilon) \text{OPT}(N')$ . Assigning small jobs greedily to a least loaded machine in step (iii) therefore results in a  $(1 + \epsilon)$ -approximate schedule for  $N'$  by Corollary 2.1.

Finally, notice that, in the whole process, only jobs that have initially been scheduled on machines  $M'$  are moved. The total size of these jobs is bounded from above by

$$2 \text{OPT}(N') |M'| \stackrel{(1)}{\leq} 4 \text{LB}(N') |M'| \leq p_j \frac{4}{\epsilon} |M'| .$$

We can therefore set  $\beta(\epsilon) := \frac{4}{\epsilon} |M'| \in 2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}$ . □

**COROLLARY 6.1** *Theorem 6.1 can be strengthened by adding the requirement that the schedules for  $N$  and  $N'$  only differ on constantly many machines.*

**PROOF.** In the proof of Theorem 6.1, the only critical step in the construction of the schedule for  $N'$  is the reassignment of small jobs. In general the number of small jobs which need to be reassigned is not constant. Moreover, in the worst case, every such small job is assigned to a different machine.

This problem can be solved by grouping small jobs into batches of size roughly  $\epsilon \text{LB}(N')$ . Each batch can be treated as one single small job. In this way we can reduce the number of small jobs which need to be reassigned to a constant. □

We can finally state the main result of this section.

**THEOREM 6.2** *There exists a  $(1 + \epsilon)$ -competitive online algorithm with constant migration factor  $\beta(\epsilon) \in 2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}$  such that the running time for incorporating a newly arrived job is constant. More precisely, the algorithm maintains and updates a data structure in time  $2^{2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}}$  in each iteration<sup>3</sup>. From this data structure a  $(1 + \epsilon)$ -approximate schedule can be obtained in time linear in the number of jobs (and not depending on  $\epsilon$ ).*

**PROOF.** Before we discuss the required data structure in detail, we first give a rough outline and intuition.

We consider again the scheduling strategy described in the proof of Theorem 6.1. In Step (i) we ignore all small jobs and construct the schedule for the large jobs  $N'_L$ . This schedule is stored as a vector  $y \in \mathbb{N}_0^{\mathbf{K}}$  (array) in our data structure. It follows from the proof of Lemma 6.2 that only a constant number of

<sup>3</sup>Here we make the common assumption that operations with numbers of polynomial size take constant time.

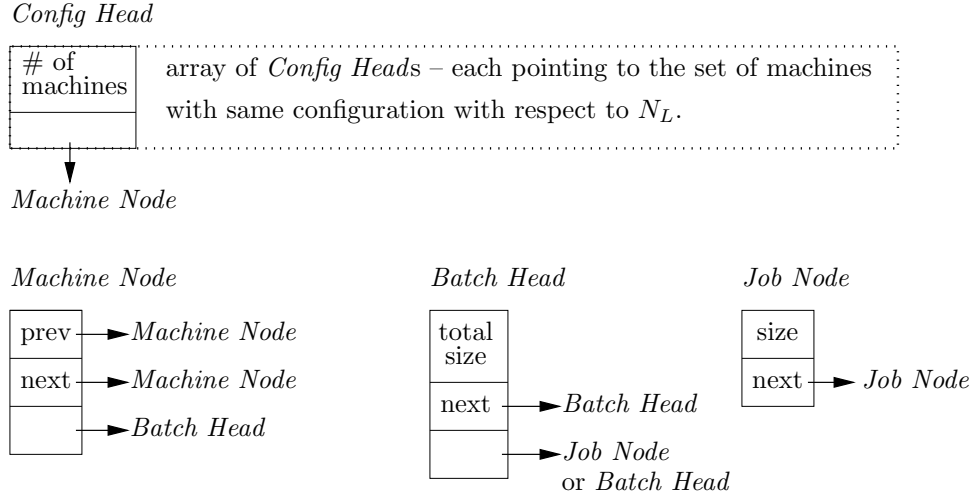


Figure 5: The data structure used to represent a schedule.

different possibilities have to be taken into consideration in order to find this schedule. The schedule for  $N'_L$  in Step (i) can thus be found in constant time. Step (ii) is trivial; nothing has to be changed on the machines in  $M \setminus M'$ .

We finally turn to Step (iii) where the small jobs that were previously assigned to  $M'$  are successively reassigned to a least loaded machine. In order to perform this step in constant time, we apply the same trick as in the proof of Corollary 6.1 and assume that small jobs are grouped into batches of size roughly  $\epsilon \text{LB}(N')$ . Each batch can be treated as one single small job such that the number of small jobs that we have to deal with in Step (iii) is constant. We finally have to argue that we can always find a least loaded machine in constant time. In fact, since we are only looking for a  $(1 + O(\epsilon))$ -approximate schedule, it is not really necessary to find a least loaded machine. It suffices to find a machine whose load differs from the minimum load by at most  $\epsilon \text{LB}(N')$ . We can therefore assume that the machines are assigned to  $O(\frac{1}{\epsilon})$  buckets according to their current load where every bucket contains machines whose load is within a range of at most  $\epsilon \text{LB}(N')$ . In Step (iii) it then suffices to always choose a machine from the non-empty bucket that contains the machines of smallest loads. This can obviously be done in  $O(\frac{1}{\epsilon})$  time.

We now give a detailed description of the data structure. We assume in the following without loss of generality that  $\epsilon < 1$ . A  $(1 + \epsilon)$ -approximate schedule for  $N$  can be represented using the following simple data structure. We assume that initially the schedule is given to us in this form. Later we show how to update it in constant time while scheduling job  $j$ . The machine configurations are represented using structures as shown in Figure 5.

There is an array of *Config Heads* of dimension  $|\mathbf{K}(4\text{LB}(N'))| \in 2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}$  (notice that  $4\text{LB}(N') \geq 2\text{OPT}(N') > (1 + \epsilon)\text{OPT}(N')$ ), one for each possible configuration  $k \in \mathbf{K}(4\text{LB}(N'))$ . Each *Config Head* corresponding to a configuration  $k$  stores the number of machines  $y_k$  scheduled according to this configuration. Moreover, each *Config Head* points to the list of machines (list of *Machine Nodes*) obeying that configuration.

Each *Machine Node* points to the list of jobs scheduled on that machine, grouped into batches. Each large job is put into a batch of its own. The small jobs on a machine are grouped into at most  $O(\frac{1}{\epsilon})$  batches of size at most  $\epsilon \text{LB}(N)$ . Thus there is only a constant number  $O(\frac{1}{\epsilon})$  of batches on any machine. Each batch has a *Batch Head* that points to the list of jobs (*Job Nodes*) in the batch. Since these jobs can recursively be grouped into smaller batches again, the list might also contain *Batch Heads*. The jobs that are scheduled on a machine are thus represented by a tree rooted at the *Machine Node* whose inner nodes are batches and whose leaves are the jobs. Since every inner node has at least two child nodes, the size of the tree is linear in the number of jobs on the machine (i.e., leaf nodes).

Finally, the machines are partitioned into  $O(\frac{1}{\epsilon})$  buckets according to their current load. Each bucket contains machines whose load lies within an interval of size at most  $\epsilon \text{LB}(N')$ .



We next discuss the details of updating the data structure when scheduling job  $j$ . If  $j$  is small (i.e.,  $p_j < \epsilon \text{LB}(N')$ ), then:

- (i) Pick any machine  $i$  from the non-empty bucket containing the least loaded machines. Since there are  $O(\frac{1}{\epsilon})$  buckets, this takes  $O(\frac{1}{\epsilon})$  time.
- (ii) Find a batch of small jobs on machine  $i$  of total size at most  $\epsilon \text{LB}(N') - p_j$  and add job  $j$  to this batch. If no such batch exists, create a new batch for job  $j$ . Consider all batches and merge small batches into larger batches of size at most  $\epsilon \text{LB}(N')$ . All this takes  $O(\frac{1}{\epsilon})$  time since there are at most  $O(\frac{1}{\epsilon})$  batches on machine  $i$ .
- (iii) Reassign machine  $i$  to a bucket according to its new load. Also this takes  $O(\frac{1}{\epsilon})$  time.

The overall time needed to assign a small job to an approximately least loaded machine is thus  $O(\frac{1}{\epsilon})$ . If the new job  $j$  is large (i.e.,  $p_j \geq \epsilon \text{LB}(N')$ ), we proceed as follows:

- (iv) Enumerate all  $2^{2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}}$  schedules (vectors)  $y' \in \mathbf{S}'$  for  $N'_L$  with  $\|y - y'\|_\infty \in 2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}$ . Choose one with minimum makespan. This takes  $2^{2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}}$  time.
- (v) The component-wise difference between  $y$  and  $y'$  specifies a subset of configurations and a non-zero number of machines that should be modified for each such configuration. For each configuration we remove the required number of machines from the front of the machine list pointed to by the respective *Config Head*. Since the total number of required machines is  $2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}$  (see Lemma 6.2), this takes  $2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}$  time.
- (vi) Remove all batches of small jobs from these machines and reschedule the remaining jobs among these machines in order to achieve the configurations indicated by  $y'$ . Since the number of jobs and batches on each machine is  $O(\frac{1}{\epsilon})$ , this takes  $2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}$  time.
- (vii) Assign the machines to their new *Config Heads* and to buckets according to their load. This takes time  $2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}$  again.
- (viii) Each of the  $2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}$  batches of small jobs is being reassigned as in the small job case discussed above. This takes  $2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}$  time.

The running time of the entire procedure is  $2^{2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}}$ .

There is one more issue that needs to be discussed. Whenever a new job arrives, the lower bound  $\text{LB}$  on the makespan of an optimal schedule might increase such that the classification of large and small jobs has to be updated. This means, that the list of configurations that are currently being used changes. Certain classes of formerly large jobs are now small and therefore neglected in the updated configurations. Notice, however, that machines belonging to the same *Config Head* before the update still belong to a common *Config Head* after the update. It is thus sufficient to update the array of *Config Heads*, potentially merging several old *Config Heads* into a new one without touching all *Machine Nodes*. This can be done in  $2^{O(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon})}$  time.

Notice that there is no need to assign the jobs that became small immediately to batches of small jobs. On each machine this is done as soon as the machine is being touched in the future. We always preserve the invariant that the number of batches of small jobs is  $O(\frac{1}{\epsilon})$  on any machine.

Due to the monotonically increasing lower bound  $\text{LB}$ , we also need to update the bucket structure from time to time by merging two buckets into one. This can obviously also be done in  $O(\frac{1}{\epsilon})$  time.

Notice finally, that from the described data structure a corresponding schedule can be obtained in linear time. For each *Machine Node* parse the tree containing the jobs scheduled on the machine. As mentioned above, the size of the tree is linear in the number of jobs contained in it. This concludes the proof.  $\square$

The migration factor  $\beta(\epsilon)$  we obtain is exponential in  $\frac{1}{\epsilon}$ . This raises the question whether a migration factor that is polynomially bounded in  $\frac{1}{\epsilon}$  is possible. We do not know of an instance proving the converse and leave the question as an open problem.

As a corollary to Theorem 6.2, we obtain the following known result; see, e.g., [17].

**COROLLARY 6.2** *There is a polynomial time approximation scheme with linear running time for the scheduling problem  $P \mid C_{\max}$ .*

**7. Maximizing the Minimum Machine Load** An alternative, yet less frequently used objective for machine scheduling is to maximize the minimum machine load. We have, however, a concrete application using this objective function that was the original motivation for our interest in bounded migration: Storage area networks (SAN) usually connect many disks of different capacity and grow over time. A convenient way to hide the complexity of a SAN is to treat it as a single big, fault tolerant disk of huge capacity and throughput [7, 26]. A simple scheme with many nice properties implements this idea if we manage to partition the SAN into several sub-servers [26] of about equal size. Mapping to the scheduling framework, the disks correspond to jobs and the sub-servers correspond to machines. Each sub-server stores the same amount of data. For example, if we have two sub-servers, each of them stores all the data to achieve a fault tolerance comparable to mirroring in ordinary RAID level 0 arrays [22]. More sub-servers allow for a more flexible tradeoff between fault tolerance, redundancy, and access granularity. In any case, the capacity of the server is determined by the *minimum* capacity of a sub-server. Moreover, it is not acceptable to completely reconfigure the system when a new disk is added to the system or when a disk fails. Rather, the user expects a “proportionate response”, i.e., if she adds a disk of  $x$  GByte she will not be astonished if the system moves data of this order of magnitude but she would complain if much more is moved. Our theoretical investigation confirms that this ‘common sense’ expectation is indeed reasonable.

We concentrate on the case without job departures (disk failures) and present a  $\frac{1}{2}$ -competitive online algorithm with migration factor 1. In order to motivate the scheduling procedure, we first discuss a negative result. The following lemma shows that it is not possible to start with an arbitrary  $\frac{1}{2}$ -approximate schedule for  $N$  and obtain a  $\frac{1}{2}$ -approximate schedule for  $N \cup \{j\}$  with constant migration factor.

LEMMA 7.1 *There is a  $\frac{1}{2}$ -approximate schedule on  $m$  machines such that upon the arrival of a new job it is not possible to obtain a  $\frac{1}{2}$ -approximate schedule with migration factor less than  $m - 2$ .*

PROOF. Consider the following  $m$ -machine schedule for  $3m - 1$  jobs of unit size. Machine 1 contains  $2m$  jobs while all remaining machines contain only one job. The optimal minimum load for this instance is 2 and the considered schedule is  $\frac{1}{2}$ -approximate. When a new job of size 1 arrives, the new optimal minimum load is 3. But to achieve minimum load greater than 1, any strategy has to move at least  $m - 2$  jobs.  $\square$

We show that the following simple strategy, which is very similar to Procedure 1, leads to a  $\frac{1}{2}$ -competitive online algorithm with migration factor 1.

PROCEDURE 5

Input: A schedule for a set of jobs  $N$  and a new job  $j \notin N$ .

Output: A schedule for  $N \cup \{j\}$ .

*Repeatedly remove jobs from the least loaded machine in any order; stop before the total size of removed jobs exceeds  $p_j$ . Assign job  $j$  to this machine. Assign the removed jobs successively to a least loaded machine.*

With respect to a given schedule a machine is called *multi-job machine* if at least two jobs are assigned to it. In the analysis of Procedure 5 we make use of a special structural property of schedules. A schedule for a set of jobs  $N$  has property (\*\*) if the load of any multi-job machine is at most twice the minimum load. For example, a schedule constructed by list scheduling in order of non-increasing processing times has always property (\*\*). We show that Procedure 5 maintains property (\*\*).

LEMMA 7.2 *Given a schedule for  $N$  satisfying property (\*\*) and a new job  $j \notin N$ , Procedure 5 outputs a schedule for  $N \cup \{j\}$  satisfying property (\*\*).*

PROOF. The proof is based on the following claim.

CLAIM: Given a schedule satisfying property (\*\*) and a new job whose size is at most the minimum machine load, assigning the new job to a least loaded machine maintains property (\*\*).

This claim is true since incorporating the new job does not decrease the minimum machine load and the load of the machine the new job is assigned to increases at most by a factor of 2.

Let Machine  $i$  be a least loaded machine in the given schedule  $S$  for  $N$ . As a consequence of the claim, reassigning the removed jobs in Procedure 5 successively to a least loaded machine maintains property (\*\*). It therefore remains to show that property (\*\*) holds in the intermediate schedule immediately after assigning job  $j$  to machine  $i$  (before assigning the removed jobs). Notice that the minimum machine load in this intermediate schedule is at least the minimum machine load of the given schedule for  $N$ . If  $p_j \leq p(S^{-1}(i))$ , then property (\*\*) obviously holds for the intermediate schedule. Otherwise, if  $p_j > p(S^{-1}(i))$ , then all jobs are removed from machine  $i$  such that machine  $i$  is not a multi-job machine in the intermediate schedule. Thus property (\*\*) holds in this case as well.  $\square$

**THEOREM 7.1** *For the objective to maximize the minimum machine load, any schedule satisfying property (\*\*) is  $\frac{1}{2}$ -approximate. In particular, scheduling every newly arriving job according to Procedure 5 is a  $\frac{1}{2}$ -competitive algorithm with migration factor 1 for this objective.*

**PROOF.** The proof is based on the following claim.

**CLAIM:** Consider an arbitrary schedule. If there is at most one job on every machine, the schedule is optimal. Otherwise, the maximum load of a multiple-job machine is an upper bound on the optimum.

The first part of the claim is clear. Consider an arbitrary schedule  $S$  with at least one machine containing at least two jobs. Let  $L$  denote the maximum load of a multiple-job machine in  $S$ . In the following we call a job *large* if its size is strictly larger than  $L$ . Notice that  $S$  schedules each large job on a machine of its own.

Consider an optimum schedule  $S^*$  and modify  $S^*$  as follows: For each machine  $i$ : If machine  $i$  contains a large job, then delete machine  $i$  and all jobs scheduled on it. The resulting schedule is denoted by  $\bar{S}^*$ . We apply exactly the same procedure to schedule  $S$  and denote the resulting schedule by  $\bar{S}$ . Notice that  $\bar{S}^*$  schedules a subset of the jobs scheduled by  $\bar{S}$  since each large job occupies a machine of its own in  $S$ . Moreover, for the same reason,  $\bar{S}^*$  uses at least as many machines as  $\bar{S}$ . Thus, the minimum machine load in  $\bar{S}^*$  is at most the maximum machine load in  $\bar{S}$  which is equal to  $L$ . In particular, there exists a machine in  $\bar{S}^*$  and thus in  $S^*$  whose load is at most  $L$ . This concludes the proof of the claim.

As an immediate consequence of the claim, any schedule satisfying property (\*\*) is  $\frac{1}{2}$ -approximate. Together with Lemma 7.2 this yields a bound of  $\frac{1}{2}$  on the competitiveness. Finally, the bound on the migration factor is clear from the description of Procedure 5.  $\square$

We close this section with an interesting open problem. Woeginger [32] presents a PTAS for the problem to maximize the minimum machine load on identical parallel machines. This raises the question whether an online PTAS with constant migration factor  $\beta(\epsilon)$  as in Section 6 can be obtained for the max-min objective function. Lemma 7.1 seems to indicate that this is more tricky than for the makespan objective. At least, it is not enough to only take care of the objective function in each iteration but also the overall structure of the schedule has to be ‘reasonable’ such that it can be easily adapted in future iterations without too much migration.

**Acknowledgments.** We thank the anonymous referees for their careful reading and their numerous suggestions which lead to an improved presentation of the paper. We would also like to thank Gerhard Woeginger for interesting discussions and helpful comments on the topic of this paper. Part of this work was done while all three authors were at Max-Planck-Institut für Informatik in Saarbrücken.

## References

- [1] F. Afrati, E. Bampis, C. Chekuri, D. Karger, C. Kenyon, S. Khanna, I. Milis, M. Queyranne, M. Skutella, C. Stein, and M. Sviridenko, *Approximation schemes for minimizing average weighted completion time with release dates*, Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science (New York City, NY), 1999, pp. 32–43.
- [2] S. Albers, *Better bounds for online scheduling*, SIAM Journal on Computing **29** (1999), 459–473.
- [3] ———, *Online algorithms: a survey*, Mathematical Programming **97** (2003), 3–26.
- [4] M. Andrews, M.X. Goemans, and L. Zhang, *Improved bounds for on-line load balancing*, Algorithmica **23** (1999), 278–301.
- [5] Y. Azar and L. Epstein, *On-line machine covering*, Journal of Algorithms **1** (1998), 67–77.

- [6] Y. Bartal, A. Fiat, H. Karloff, and R. Vohra, *New algorithms for an ancient scheduling problem*, Journal of Computer and System Sciences **51** (1995), 359–366.
- [7] A. Brinkmann, K. Salzwedel, and C. Scheideler, *Compact, adaptive placement schemes for non-uniform requirements*, 14th ACM Symposium on Parallel Algorithms and Architectures, 2002, pp. 53–62.
- [8] B. Chen, A. van Vliet, and G. J. Woeginger, *Lower bounds for randomized online scheduling*, Information Processing Letters **51** (1994), 219–222.
- [9] E. G. Coffman Jr., M. R. Garey, and D. S. Johnson, *An application of bin-packing to multiprocessor scheduling*, SIAM Journal on Computing **7** (1978), 1–17.
- [10] L. Epstein and A. Levin, *A robust APTAS for the classical bin packing problem*, Automata, Languages and Programming (M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, eds.), Lecture Notes in Computer Science, vol. 4051, Springer, Berlin, 2006, pp. 214–225.
- [11] R. Fleischer and M. Wahl, *Online scheduling revisited*, Journal of Scheduling **3** (2000), 343–353.
- [12] D. K. Friesen, *Tighter bounds for the multifit processor scheduling algorithm*, SIAM Journal on Computing **13** (1984), 170–181.
- [13] M. R. Garey and D. S. Johnson, *Strong np-completeness results: Motivation, examples and implications*, Journal of the ACM **25** (1978), 499–508.
- [14] R. L. Graham, *Bounds for certain multiprocessing anomalies*, Bell System Technical Journal **45** (1966), 1563–1581.
- [15] \_\_\_\_\_, *Bounds on multiprocessing timing anomalies*, SIAM Journal on Applied Mathematics **17** (1969), 263–269.
- [16] R. L. Graham, E. L. Lawler, J. K. Lenstra, and A. H. G. Rinnooy Kan, *Optimization and approximation in deterministic sequencing and scheduling: A survey*, Annals of Discrete Mathematics **5** (1979), 287–326.
- [17] D. S. Hochbaum, *Various notions of approximation: Good, better, best, and more*, Approximation Algorithms for NP-Hard Problems (D. S. Hochbaum, ed.), Thomson, 1996, pp. 346–398.
- [18] D. S. Hochbaum and D. B. Shmoys, *Using dual approximation algorithms for scheduling problems: Theoretical and practical results*, Journal of the ACM **34** (1987), 144–162.
- [19] D. R. Karger, S. J. Phillips, and E. Torng, *A better algorithm for an ancient scheduling problem*, Journal of Algorithms **20** (1996), 400–430.
- [20] M. A. Langston, *Processor scheduling with improved heuristic algorithms*, Ph.D. thesis, Texas A&M University, 1981.
- [21] H. W. Lenstra, *Integer programming with a fixed number of variables*, Mathematics of Operations Research **8** (1983), 538–548.
- [22] D. Patterson, G. Gibson, and R. Katz, *A case for redundant arrays of inexpensive disks (RAID)*, Proceedings of ACM SIGMOD’88 (1988), 109–116.
- [23] J. F. Rudin III, *Improved bounds for the on-line scheduling problem*, Ph.D. thesis, The University of Texas at Dallas, 2001.
- [24] J. F. Rudin III and R. Chandrasekaran, *Improved bounds for the online scheduling problem*, SIAM Journal on Computing **32** (2003), 717–735.
- [25] S. Sahni, *Algorithms for scheduling independent tasks*, Journal of the ACM **23** (1976), 116–127.
- [26] P. Sanders, *Algorithms for scalable storage servers*, SOFSEM 2004: Theory and Practice of Computer Science, vol. 2932, 2004, pp. 82–101.
- [27] P. Sanders, N. Sivadasan, and M. Skutella, *Online scheduling with bounded migration*, Automata, Languages and Programming (J. Diaz, J. Karhumäki, A. Lepistö, and D. Sannella, eds.), Lecture Notes in Computer Science, vol. 3142, Springer, Berlin, 2004, pp. 1111–1122.
- [28] A. Schrijver, *Theory of Linear and Integer Programming*, John Wiley & Sons, Chichester, 1986.
- [29] J. Sgall, *A lower bound for randomized on-line multiprocessor scheduling*, Information Processing Letters **63** (1997), no. 1, 51–55.
- [30] J. Sgall, *On-line scheduling — a survey*, Online Algorithms: The State of the Art (A. Fiat and G. J. Woeginger, eds.), Lecture Notes in Computer Science, vol. 1442, Springer, Berlin, 1998, pp. 196–231.

- [31] J. Westbrook, *Load balancing for response time*, J. Algorithms **35** (2000), no. 1, 1–16.
- [32] G. J. Woeginger, *A polynomial-time approximation scheme for maximizing the minimum machine completion time*, Operations Research Letters **20** (1997), 149–154.