Approximation Schemes for Scheduling on Uniformly Related and Identical Parallel Machines

Leah Epstein* Jiří Sgall†

March 27, 2003

Abstract

We give a polynomial approximation scheme for the problem of scheduling on uniformly related parallel machines for a large class of objective functions that depend only on the machine completion times, including minimizing the l_p norm of the vector of completion times. This generalizes and simplifies many previous results in this area.

Keywords: Scheduling, approximation algorithms, parallel machines, machine completion time, polynomial time approximation scheme.

1 Introduction

Scheduling is one of fundamental areas of combinatorial optimization. Most multiprocessor scheduling problems are known to be hard to solve optimally (NP-hard, see below). Thus the research focuses on giving efficient approximation algorithms that produce a solution close to the optimal one. Ideally, one hopes to obtain a family of polynomial algorithms such that for any given $\varepsilon > 0$ the corresponding algorithm is guaranteed to produce a solution with a cost within a factor of $(1 + \varepsilon)$ of the optimum cost; such a family is called a polynomial approximation scheme.

A polynomial scheme for the basic problem of minimization of the total length of the schedule (the makespan) on identical machines was given by Hochbaum and Shmoys [12]. This technique was later generalized to a number of scheduling (and other) problems.

Our paper gives a unified and comprehensive treatment of various special cases studied before and their further generalization. We believe that our presentation is easy to read,

*E-mail lea@idc.ac.il, School of Computer Science, The Interdisciplinary Center, Herzliya, Israel. This work carried out while the author was at Tel-Aviv University.

†E-mail sgall@math.cas.cz, http://www.math.cas.cz/~sgall, Mathematical Inst., AS CR, Žitná 25, CZ-11567 Praha 1, Czech Republic and Institute for Computer Science, Prague (project LN00A056 of MŠMT CR). Partially supported by grants 201/97/P038 and 201/01/1195 of GA ČR, and cooperative research grant KONTAKT-ME476/CCR-9988360-001 from the NSF and MŠMT CR.

perhaps even easier than the specialized results published before, since the general treatment let us focus on the important issues. Our algorithms can be applied to many special scheduling problems, either directly, or with minor modification. This is demonstrated on the example of scheduling with rejection in Section 7.

We are given n jobs with processing times p_j, $j \in J = \{1, \ldots, n\}$ and m machines M_1, M_2, ..., M_m, with speeds s_1, s_2, ..., s_m. A schedule is an assignment of the n jobs to the m machines. Given a schedule, for $1 \leq i \leq m$, T_i denotes the weight of machine M_i which is the total processing time of all jobs assigned to it, and C_i denotes the completion time of M_i, which is T_i/s_i. (Each job is assigned to exactly one machine, i.e., we do not allow preemption.)

Our objective is, for some fixed function $f : [0, +\infty) \to [0, +\infty)$, one of the following:

(I) minimize	$\sum_{i=1}^{m} f(C_i)$,	**(III)** maximize	$\sum_{i=1}^{m} f(C_i)$, or
(II) minimize	$\max_{i=1}^{m} f(C_i)$,	**(IV)** maximize	$\min_{i=1}^{m} f(C_i)$.

Most of such problems are NP-hard, see [9, 14]. Thus we are interested in approximation algorithms. Recall that a *polynomial time approximation scheme (PTAS)* is a family of polynomial time algorithms over $\varepsilon > 0$ such that for every ε and every instance of the problem, the corresponding algorithm outputs a solution whose value is within a factor of $(1 + \varepsilon)$ of the optimum value [9]. Since even the simple cases of our problem are NP-hard in the strong sense [9], we cannot hope to get a fully polynomial approximation scheme (then we would require the running time to depend polynomially also on ε, and strong NP-hardness of the problems implies that such an algorithm would show that $P = NP$).

We give a PTAS for scheduling on uniformly related machines for a rather general set of functions f, covering many natural functions studied before. Let us give some examples covered by our results and references to previous work.

Example 1. Problem (II) with $f(x) = x$ is the basic problem of minimizing the maximal completion time (makespan). It was studied for identical machines in [10, 11, 15, 12]; the last paper gives a PTAS. Finally, Hochbaum and Shmoys [13] gave a PTAS for uniformly related parallel machines. Thus our result can be seen as a generalization of that result. In fact our paper is based on their techniques, perhaps somewhat simplified.

Example 2. Problem (I) with $f(x) = x^p$ for any $p > 1$. This is equivalent to minimizing the l_p norm of the vector $(C_1, \ldots, C_m)$, i.e., $(\sum C_i^p)^{1/p}$. For $p = 2$ and identical machines this problem was studied in [5, 4], motivated by storage allocation problems. For a general p a PTAS for identical machines was given in [1]. For related machines this problem was not studied before.

Note that for $p \leq 1$ the minimization problem is trivial: it is optimal to schedule all jobs on the fastest machine (choose one arbitrarily if there are more of them). In this case a more interesting variant is the maximization version, i.e., the problem (III).

Example 3. Problem (IV) with $f(x) = x$. In this problem the goal is to maximize the time when all the machines are running; this corresponds to keeping all parts of some system alive as long as possible. This problem on identical machines was studied in [8, 6], and [16] gave a PTAS. For uniformly related machines a PTAS was given in [2].

Example 4. Scheduling with rejection. In this problem each job has associated certain penalty. The schedule is allowed to schedule only some subset of jobs, and the goal is to minimize the maximal completion time plus the total penalty of all rejected jobs. This problem does not conform exactly to any of the categories above, nevertheless our scheme can be extended to work for it as well. This problem was studied in [3] where also a PTAS for the case of identical machines is given. For related machines this problem was not studied.

Our paper is directly motivated by Alon et al. [1], who proved similar general results for scheduling on identical machines. We generalize it to uniformly related machines and a similar set of functions f. Even for identical machines, our result is stronger than that of [1] since in that case we allow a more general class of functions f. In Section 6 we also disprove a conjecture of [1] concerning the class of functions f allowing PTAS for identical machines.

The basic idea is to round the size of all jobs to a constant number of different sizes, to solve the rounded instance exactly, and then re-construct an almost optimal schedule for the original instance. This rounding technique traces back to Hochbaum and Shmoys [12]. We also use an important improvement from [1]: the small jobs are clustered into blocks of jobs of small but non-negligible size. The final ingredient is that of [13, 2]: the rounding factor is different for each machine, increasing with the weight of the machine (i.e., the total processing time of jobs assigned to it). Such rounding is possible if we assign the jobs to the machines in the order of non-decreasing weight. This is easy to do for identical machines. For uniformly related machines we prove, under a reasonable condition on the function f, that in a good solution the weight of machines increases (or decreases) with their speed, which fixes the order of machines.

2 Our assumptions and results

Now we state our assumptions on the function f. Let us note at the beginning that the typical functions used in scheduling problems satisfy all of them.

The first condition says that to approximate the contribution of a machine up to a small multiplicative error, it is sufficient to approximate the completion time of the machine up to a small multiplicative error. This condition is from Alon et al. [1], and it is essential for all our results.

$$\begin{aligned}(\mathrm{F}*):\quad & (\forall \varepsilon > 0)(\exists \delta > 0)(\forall x, y \geq 0) \\ & (|y-x| \leq \delta x \rightarrow |f(y) - f(x)| \leq \varepsilon f(x)).\end{aligned}$$

If we transform f so that both axes are given a logarithmic scale, the condition naturally translates into uniform continuity, as $|y-x| \leq \delta x$ iff $(1-\delta)x \leq y \leq (1+\delta)x$ iff $\ln(1-\delta) + \ln x \leq \ln y \leq \ln(1+\delta) + \ln x$, and similarly for $f(x)$. More formally, (F$*$) is equivalent to the following statement:

$$(\mathrm{F}**):\quad \text{The function } h_f : (-\infty, +\infty) \rightarrow (-\infty, +\infty)$$

defined by $h_f(z) = \ln f(e^z)$
is defined everywhere and uniformly continuous.

For the typical case of convex functions, this implies that f grows at most polynomially, see Section 6 for more discussion of the condition.

We also need a condition that would guarantee that in the optimal schedule, the weights of the non-empty machines are monotone. These conditions are different for the cases when the objective is maximize or minimize $\sum f(C_i)$, and for the cases of min-max or max-min objectives.

Recall that a function g is convex iff for every $x \le y$ and $0 \le \Delta \le y - x$, $f(x + \Delta) + f(y - \Delta) \le f(x) + f(y)$. For the cases of min-sum and max-sum we require this condition:

$$(\mathrm{G}*): \quad \text{The function } g_f : (-\infty, +\infty) \to [0, +\infty) \text{ defined by } g_f(z) = f(e^z) \text{ is convex.}$$

Note that $f(x) = g_f(\ln x)$. Thus, the condition says that the function f is convex, if plotted in a graph with a logarithmic scale on the x-axis and a linear scale on the y-axis. This is true for example for any non-decreasing convex function f. However, g_f is convex, e.g., even for $f(x) = \ln(1 + x)$. On the other hand, the condition (G$*$) implies that f is either non-decreasing or it is unbounded for x approaching 0.

For the case of min-max or max-min we require that the function f is bimodal on $(0, +\infty)$, i.e., there exists an x_0 such that f is monotone (non-decreasing or non-increasing) both on $(0, x_0]$ and $[x_0, +\infty)$. (E.g., $(x - 1)^2 + 1$ is bimodal, decreasing on $(0, 1]$ and increasing on $[1, +\infty)$.) This includes all convex functions as well as all non-decreasing functions.

Note that none of the conditions above puts any constraints on $f(0)$. The function f can be even discontinuous at 0, which means that the cost of an empty machine can be very different from a cost of a machine with a single small job assigned to it.

Last, we need the function f to be computable in the following sense: for any $\varepsilon > 0$ there exists an algorithm that on any rational x outputs a value between $(1 - \varepsilon)f(x)$ and $(1 + \varepsilon)f(x)$, in time polynomial in the size of x. To simplify the presentation, we assume in the proofs that f is computable exactly; choosing smaller ε and computing the approximations instead of the exact values always works. Typical functions f that we want to use are computable exactly, but for example if $f(x) = x^p$ for non-integral p then we can only approximate it.

Our main results are:

Theorem 2.1 *Let f be a non-negative computable function satisfying the conditions* (F$*$) *and* (G$*$). *Then the scheduling problems of minimizing and maximizing $\sum f(C_i)$ on uniformly related machines both possess a PTAS.*

Theorem 2.2 *Let f be a non-negative computable function satisfying the condition* (F$*$) *and bimodal on* $(0, +\infty)$. *Then the scheduling problems of minimizing* $\max f(C_i)$ *and of maximizing* $\min f(C_i)$ *on uniformly related machines both possess a PTAS.*

Theorem 2.3 *Let f be a non-negative computable function satisfying the condition* (F$*$). *Then all the scheduling problems of minimizing or maximizing $\sum f(C_i)$, of minimizing* $\max f(C_i)$*, and of maximizing* $\min f(C_i)$ *on identical machines possess a PTAS.*

All our PTAS are running in time $O(n^c p(|I|))$, where c is a constant depending on the function f and the desired precision ε, p is the polynomial bounding the time of the computation of f, and $|I|$ is the size of the input instance. Thus the time is polynomial, but the exponent in the polynomial depends on f and ε. This should be contrasted with Alon et al [1], where for the case of identical machines they are able to achieve linear time (i.e., the exponential dependence on ε is only hidden in the constant) using integer programming in fixed dimension. It is an open problem if such an improvement is also possible for related machines; this is open even for the case of minimizing the makespan (Example 1 above).

3 Ordering of the machines

The following lemma implies that, depending on the type of the problem and f, we can order the machines in either non-decreasing or non-increasing order of speeds and then consider only the schedules in which the weights of the machines are non-decreasing (possibly with the exception of the empty machines). It shows why the conditions (G$*$) and bimodality are important for the respective problems; it is the only place where the conditions are used. Note that for identical machines the corresponding statements hold trivially without any condition.

Lemma 3.1 *Let the machines be ordered so that the speeds s_i are non-decreasing.*

(i) *Under the same assumptions as in Theorem 2.1, there exists a schedule with minimal (maximal, resp.) $\sum f(C_i)$ in which the non-zero weights of the machines are monotone non-decreasing (monotone non-increasing, resp.). I.e., for any $1 \le i < j \le m$ such that $T_i, T_j > 0$, we have $T_i \le T_j$ ($T_i \ge T_j$, resp.).*

(ii) *Under the same assumptions as in Theorem 2.2, there exist schedules both with minimal* $\max f(C_i)$ *and with maximal* $\min f(C_i)$ *in which either (a) the non-zero weights of the machines are monotone non-decreasing or (b) for some $1 \le k < m$, the k smallest jobs are assigned to the k fastest machines, one job per machine, in non-decreasing order, and the remaining machines have non-zero and non-decreasing weights.*

Proof: (i) We prove the lemma for minimization of $\sum f(C_i)$; the case of maximization is similar (with the order reversed). We prove that if $T_i > T_j$ for some $i < j$ then switching the assigned jobs between M_i and M_j leads to at least as good schedule. This is sufficient, since given any optimal schedule we can obtain an optimal schedule with ordered weights by at most $m-1$ such transpositions.

Denote $s = s_j/s_i$, $\Delta = \ln s$, $X = \ln C_j = \ln(T_j/s_j)$, and $Y = \ln C_i = \ln(T_i/s_i)$. From the assumptions we have $X < Y$ and $0 \le \Delta < Y - X$. The difference of the original cost and the cost after the transposition is $f(C_j) + f(C_i) - f(sC_j) - f(C_i/s) = g_f(X) + g_f(Y) - g_f(X + \Delta) - g_f(Y - \Delta) \geq 0$, using convexity of g_f. Thus the transposed schedule is at least as good as the original one.

(ii) We prove the lemma for minimization of $\max f(C_i)$; the case of maximization of $\min f(C_i)$ is symmetric. We have two subcases.

First suppose that we are minimizing $\max f(C_i)$ and f is bimodal, first non-increasing then non-decreasing. We prove that (a) is satisfied by showing that if $T_i > T_j$ for $i < j$, then switching the assigned jobs between M_i and M_j can only improve the schedule. Let $s = s_j/s_i \geq 1$. By the assumptions we have $C_j \leq sC_j \leq C_i$ and $C_j \leq C_i/s \leq C_i$. By bimodality of f we have $\max\{f(sC_j), f(C_i/s)\} \leq \max\{f(C_j), f(C_i)\}$, hence the transposed schedule can only be better.

Now suppose that we are minimizing $\max f(C_i)$ with f first non-decreasing and then non-increasing. Choose x_0 such that f is non-decreasing on $(0, x_0)$ and then non-increasing and consider an optimal schedule. If there exist two machines M_i with $T_i = 0$ and M_j with $C_j > x_0$, we can get at least as good solution by scheduling all the jobs on M_j and (a) is satisfied. If there exist no nonempty machines M_i and M_j with $0 < C_i < x_0 < C_j$ then we are done by the previous case, as we may restrict our attention to one of the intervals where f is monotone. If machines M_i and M_j are such that $0 < C_i < x_0 < C_j$ then we may assume that $i > j$ and M_i contains a single job smaller than all jobs on M_j: otherwise we obtain at least as good a schedule by switching all the jobs on the two machines, moving a job from M_i to M_j, or switching a single job from M_i for a smaller job on M_j. Thus, for some $1 \leq k < m$, the k fastest machines contain the k smallest jobs, one job per machine. As in the previous case, it follows that these k machines have non-decreasing weights, and the remaining machines also have non-decreasing (and non-zero) weights, thus (b) is satisfied. ■

Note that even though the last case of the proof is longer than the other two, typically it is the case when the scheduling problem is trivial. We include it only to have a simpler condition in the final result.

4 Preliminaries and definitions

Let $\delta > 0$ and λ be such that $\lambda = 1/\delta$ is an even integer; we will choose it later. The meaning of δ is the (relative) rounding precision.

Given w, either 0 or an integral power of two, intuitively the order of magnitude, we want to represent a set of jobs with processing times not larger than w as follows. Each job of size more than δw is replaced by a job with processing time rounded to the next higher multiple of $\delta^2 w$, while the remaining small jobs are replaced by an appropriate number of jobs with processing time δw. Now it is sufficient to remember the number n_i of modified jobs of processing time $i\delta^2 w$, for each i, $\lambda \leq i \leq \lambda^2$. Such a vector together with w is called a configuration.

In our approximation scheme, we proceed machine by machine, and use this representation to remember the set of all jobs scheduled so far. We always use the least possible

w (principal configurations below), to make the additive error caused by the small jobs relatively small; for this we also need the condition of non-decreasing weight of machines.

There are some technical difficulties in this approach. First, we need to represent the empty machines exactly, as otherwise it would be impossible to handle the additive error. Second, as we proceed the order of magnitude w increases and if we rounded the sizes of jobs repeatedly, the error could accumulate. To avoid this, we round the sizes of jobs in advance, always with respect to the largest w such that the job is larger than δw; this is implicit in the use of rounding function r below (and corrects an error from the conference version [7]). Since the values of w are powers of two, this ensures that the size is the necessary integral multiple also for smaller relevant values of w. Still, the configurations need to be carefully rescaled when changing w. These problems make the following definitions somewhat more technical than for the case of identical machines where no rescaling is needed.

Definition 4.1 *Let $A \subseteq J$ be a set of jobs.*

- *A rounding function $r(p)$ for $p > 0$ is defined as follows. Let w be the largest power of two such that $p > \delta w$. Let i be the smallest integer such that $p \le i\delta^2 w$. Then $r(p) = i\delta^2 w$. We define $r(0) = 0$. (Note that whenever w' is an integer power of two and $\delta w' < p \le w'$ holds, $r(p)$ is an integral multiple of $\delta^2 w'$.)*

- *A* configuration *is a pair $\alpha = (w, (n_\lambda, n_{\lambda+1}, \dots, n_{\lambda^2}))$, where $w = 0$ or $w = 2^i$ for some integer i (possibly negative) and $\vec{n}$ is a vector of nonnegative integers.*

- *The* weight of a set of jobs *is $W(A) = \sum_{j \in A} p_j$.*

  The rounded weight of a set of jobs *is $Wr(A) = \sum_{j \in A} r(p_j)$.*

  The weight of a configuration *$(w, \vec{n})$ is defined by $W(w, \vec{n}) = \sum_{i=\lambda}^{\lambda^2} n_i i\delta^2 w$.*

- *Given a set of jobs A and a weight w, a* set of small jobs *is $A(w) = \{j \in A \mid p_j \le \delta w\}$.*

- *A configuration $(w, \vec{n})$* represents *A if*

  (i) *no job $j \in A$ has processing time $p_j > w$,*

  (ii) *for any i, $\lambda < i \le \lambda^2$, n_i equals the number of jobs $j \in A$ with $r(p_j) = i\delta^2 w$, and*

  (iii) *$n_\lambda \in \{\lfloor Wr(A(w))/(\delta w) \rfloor, \lceil Wr(A(w))/(\delta w) \rceil\}$; this is equivalent to $|n_\lambda \delta w - Wr(A(w))| < \delta w$.*

- *The* principal configuration of *A is a configuration $\alpha(A) = (w, \vec{n})$ with the smallest w which represents A; of the (at most) two such configurations, it is the one with the larger n_λ, i.e., $n_\lambda = \lceil Wr(A(w))/(\delta w) \rceil$.*

- *A a configuration $\alpha(A) = (w, \vec{n})$ is a* principal configuration *if it is the principal configuration of any $A \subseteq J$, i.e., if there exists $A \subseteq J$ such that $(w, \vec{n}) = \alpha(A)$.*

- *The* scaled configuration *for* $(w, \vec{n})$ *and* $w' \geq w$ *is defined as a vector* $\mathrm{scale}_{w \to w'}(n) = \vec{n}'$ *such that* $(w', \vec{n}')$ *represents the set* K *containing exactly* n_i *jobs with processing time equal to* $i\delta^2 w$*, for each* $i = \lambda, \ldots, \lambda^2$*, and no other jobs; of the (at most) two such configurations choose the one satisfying* $|Wr(K(w')) - n'_\lambda \delta w'| \leq \delta w'/2$*, breaking ties arbitrarily, if both configurations satisfy the inequality. (Note that for* $w' = w$*,* $\mathrm{scale}_{w \to w}(\vec{n}) = \vec{n}$*.)*

The definition implies that for each A and sufficiently large w there are exactly one or two configurations representing it, and they can be computed efficiently. Similarly, the principal configurations of a set of jobs and scaled configurations can be computed efficiently.

The crucial point of our algorithm is that instead of enumerating all the (exponentially many) sets of jobs it is sufficient to enumerate principal configurations and bound their number, which is done in the next lemma. A single principle configuration may represent many different sets of jobs that are equivalent for us. It is easy to see that a configuration $(w, \vec{n}')$ is principal if and only if (i) $\vec{n}' \leq \vec{n}$ coordinatewise where $(w, \vec{n}')$ is the principal representation of $J(w) = \{j \in J \mid p_j \leq w\}$ and (ii) $n'_i > 0$ for some $i > \lambda^2/2$. The first condition guarantees that the set of all jobs contains sufficiently many jobs of each size. The second condition guarantees that any set represented by $(w, \vec{n}')$ contains a set of size larger than $(\lambda^2/2)\delta^2 w = w/2$ and thus the chosen w is the smallest one that can represent it. (Here we use the assumption that λ is even.)

Lemma 4.2 *The number of principal configurations is at most* $(n+1)^{\lambda^2}$ *and they can be enumerated efficiently.*

Proof: There are at most $n+1$ possible values of w in the principal configurations: 0 and p_j's rounded up to a power of two. To enumerate all principal configurations with a given w, find a representation $(w, \vec{n})$ of $J(w) = \{j \in J \mid p_j \leq w\}$ and enumerate all the vectors $\vec{n}'$ bounded by $\vec{0} \leq \vec{n}' \leq \vec{n}$ (coordinatewise), such that $n'_i > 0$ for some $i > \lambda^2/2$. For each i, there are at most $n+1$ possible values. ∎

The next lemma proves that scaling a configuration which represents a given set produces a representing configuration as well (this property is the reason why we need to consider two representations of a set of job with any given weight).

Lemma 4.3 *If* A *is represented by* $(w, \vec{n})$ *then it is also represented by* $(w', \mathrm{scale}_{w \to w'}(\vec{n}))$ *for any* $w' \geq w$*.*

Proof: If $w = 0$ or $w' = w$ then the statement is trivial. Otherwise, condition (i) of the definition of representation is satisfied since $w' > w$. Condition (ii) is satisfied since all the jobs with $p_j > \delta w'$ are accounted for in the definition of $\mathrm{scale}_{w \to w'}$ exactly. To verify (iii), we assume, w.l.o.g., that $A = A(w')$, i.e., A contains only jobs with processing time at most $\delta w'$. Let K be the set of jobs from the definition of $\mathrm{scale}_{w \to w'}$, and let $\vec{n}' = \mathrm{scale}_{w \to w'}(A)$. Since A is represented by $(w, \vec{n})$, the definition of K implies $|Wr(A) -$

$Wr(K)| = |Wr(A(w)) - Wr(K(w))| < \delta w \leq \delta w'/2$. From the definition of scale and $K = K(w')$ we have $|Wr(K) - n'_\lambda \delta w'| \leq \delta w'/2$. By summing these two bounds we obtain $|Wr(A) - n'_\lambda \delta w'| < \delta w'$, and $(w', \vec{n}')$ represents A. ∎

The next lemma bounds the error caused by approximation. The set $B - A$ corresponds to jobs assigned to a machine M_i if A is the set of jobs assigned to previous machines and B are all jobs assigned including jobs assigned to M_i.

Lemma 4.4 *Let $A \subset B \subseteq J$ be sets of jobs and $(w, \vec{n})$ and $(w, \vec{n}')$ be any two configurations representing A and B, respectively. Then*

$$\begin{aligned} W(B - A) - 2\delta w \ &\leq \ Wr(B - A) - 2\delta w \ < \ W(w, \vec{n}' - \vec{n}) \\ &< \ Wr(B - A) + 2\delta w \ \leq \ (1 + \delta)W(B - A) + 2\delta w \end{aligned}$$

Proof: By the definition of r, for any job j, $p_j \leq r(p_j) < (1 + \delta)p_j$. Summing over all jobs $j \in B - A$ we get $W(B - A) \leq Wr(B - A) < (1 + \delta)W(B - A)$.

The contributions of any job not in $B(w)$ to $Wr(B - A)$ and to $W(w, \vec{n}' - \vec{n})$ are equal by the definition of a representing configuration. By the bounds on n_λ and n'_λ from the definition of a representing configuration, we have $|Wr(B - A) - W(w, \vec{n}' - \vec{n})| \leq |Wr(A(w)) - n_\lambda \delta w| + |Wr(B(w)) - n'_\lambda \delta w)| < 2\delta w$. (Note that $w > 0$ since $B \neq \emptyset$.) The lemma follows by combining the proven inequalities. ∎

The last lemma we need allows to find a set of jobs to be assigned to a machine if we are given the set of jobs already used and the configuration we are supposed to reach after assigning to this machine. In the degenerate case of $A = \emptyset$ (i.e., the first machine), we are simply given a principal configuration and need to find some corresponding set of jobs. In general, we need to extend this to the situation when we are also given a subset A of the set of jobs we are seeking.

Lemma 4.5 *Let $(w, \vec{n})$ be a configuration representing $A \subseteq J$. Let $w' \geq w$ and $\vec{n}'' = \mathrm{scale}_{w \to w'}(n)$. Let $(w', \vec{n}')$ be a principal configuration satisfying $\vec{n}' \geq \vec{n}''$, Then there exists a set of jobs B represented by $(w', \vec{n}')$ such that $A \subseteq B \subseteq J$, and it can be constructed in linear time.*

Proof: Construct B from A as follows. For each i, $\lambda < i \leq \lambda^2$, add $n'_i - n''_i$ jobs with processing time satisfying $r(p_j) = i\delta^2 w$; since $(w', \vec{n}')$ is a principal configuration we are guaranteed that a sufficient number of such jobs exists. Finally, add jobs with $p_j \leq \delta w'$ one by one until their weight is strictly larger than $(n'_\lambda - 1)\delta w'$. We have sufficiently many of the small jobs as well since $(w', \vec{n}')$ is principal. If any small jobs are added, the final weight is at most $n'_\lambda \delta w'$, as each added job increases the weight by at most $\delta w'$. If no small job is added, A is represented by $(w', \vec{n}'')$ by Lemma 4.3, thus $Wr(A(w')) < (n''_\lambda + 1)\delta w' \leq (n'_\lambda + 1)\delta w'$. In both cases the set B is represented by $(w', \vec{n}')$. ∎

5 The approximation scheme

Given an $\varepsilon \in (0, 1]$, we choose $\delta \leq 1/12$ using (F$*$) so that $\lambda = 1/\delta$ is an even integer and

$$(\forall x, y \geq 0)(|y - x| \leq 8\delta x \rightarrow |f(y) - f(x)| \leq \frac{\varepsilon}{3} f(x)).$$

Definition 5.1 *We define the graph G of configurations as follows.*

The vertices of G are $(i, \alpha(A))$, for any $1 \leq i < m$ and any $A \subseteq J$, the source vertex $(0, \alpha(\emptyset))$, and the target vertex $(m, \alpha(J))$.

For any i, $1 \leq i \leq m$, and any principal configurations $(w, \vec{n})$ and $(w', \vec{n}')$ with $w' \geq w$, let $\vec{n}'' = \mathrm{scale}_{w \to w'}(\vec{n})$. There is an edge from $(i-1, (w, \vec{n}))$ to $(i, (w', \vec{n}'))$ iff either $(w, \vec{n}) = (w', \vec{n}')$, or $\vec{n}'' \leq \vec{n}'$ and $W(w', \vec{n}' - \vec{n}'') \geq w'/3$. The cost of this edge is $f(W(w', \vec{n}' - \vec{n}'')/s_i)$. There are no other edges.

Definition 5.2 *Let $J_1, \ldots, J_m$ be a schedule assigning jobs in J_i to machine M_i. A sequence $\{(i, (w_i, \vec{n}_i))\}_{i=0}^{m}$ of vertices of G represents (is a principal representation of, resp.) the assignment if, for each $1 \leq i \leq m$, $(w_i, \vec{n}_i)$ represents (is a principal representation of, resp.) $\bigcup_{i'=1}^{i} J_{i'}$.*

The approximation scheme performs the following steps:

(1) Order the machines with speeds either non-decreasing or non-increasing, according to the type of the problem and f so that by Lemma 3.1 there exists an optimal schedule with non-decreasing non-zero weights of the machines. (See step (5) for the exceptional case when no such optimal schedule exists.)

(2) Construct the graph G.

(3) Find an optimal path in G from source $(0, \alpha(\emptyset))$ to $(m, \alpha(J))$. The cost of the path is defined as the sum, maximum or minimum of the costs of the edges used, and an optimal path is one with the cost minimized or maximized, as specified by the problem.

(4) Output an assignment represented by the optimal path constructed as follows: Whenever the path contains an edge of the form $((i - 1, (w, \vec{n})), (i, (w, \vec{n})))$, put $J_i = \emptyset$. For every other edge, apply Lemma 4.5, starting from the beginning of the path.

(5) In the case of minimizing the maximum or maximizing the minimum for a bimodal function f which is first non-decreasing and then non-increasing, repeat the previous steps also for $k = 1, \ldots, m-1$ so that the k smallest jobs are assigned to the k fastest machines and these jobs and machines are removed for execution of the previous steps. Compare all m resulting schedules and choose the best one.

Lemma 4.2 shows that we can construct the vertices of G in time polynomial in n. Computing the edges of G and their costs is also efficient. Since the graph G is layered, finding an optimal path takes linear time in the size of G. Given a path in a graph, finding a corresponding assignment is also fast. Hence the complexity of our PTAS is as claimed.

Lemma 5.3

(i) *Let $\{J_i\}$ be an assignment with non-decreasing weights of the machines with non-zero weights (cf. Lemma 3.1). Then its principal representation $\{(i,(w_i,\vec{n}_i))\}_{i=0}^m$ is a path in G.*

(ii) *Let $\{(i,(w_i,\vec{n}_i))\}_{i=0}^m$ be a path in G representing an assignment $\{J_i\}$ such that if $(w_{i-1},\vec{n}_{i-1}) = (w_i,\vec{n}_i)$ then $J_i=\emptyset$. Let C be the cost of a schedule given by the assignment, and let $C^\#$ be the cost of the representation as a path in the graph. Then $|C-C^\#| \le \varepsilon C/3$.*

Proof: Let $\vec{n}'_{i-1} = \mathrm{scale}_{w_{i-1}\to w_i}(\vec{n}_{i-1})$, for $1 \le i \le m$. By Lemma 4.3, $(w_i,\vec{n}'_{i-1})$ represents the set $\bigcup_{i'=1}^{i-1} J_{i'}$.

(i) First note that we really obtain vertices of G, as $(w_0,\vec{n}_0)=\alpha(\emptyset)$ and $(w_m,\vec{n}_m) = \alpha(J)$. For any $i=1,\dots,m$, if $J_i=\emptyset$ then $(w_{i-1},\vec{n}_{i-1}) = (w_i,\vec{n}_i)$ and by definition $((i-1,(w_{i-1},\vec{n}_{i-1})),(i,(w_i,\vec{n}_i)))$ is an edge of G. Otherwise, since $(w_i,\vec{n}'_{i-1})$ represents the set $\bigcup_{i'=1}^{i-1} J_{i'}$ and $(w_i,\vec{n}_i)$ is the principal representation of its superset $\bigcup_{i'=1}^{i} J_{i'}$, we have $\vec{n}'_{i-1} \le \vec{n}_i$. It remains to verify that $W(w_i,\vec{n}_i-\vec{n}'_{i-1}) \ge w_i/3$. Since $(w_i,\vec{n}_i)$ is a principal representation of $\bigcup_{i'=1}^{i} J_{i'}$, this set contains a job with $p_j > w_i/2$. Since the assignment has non-decreasing weights, it follows that $W(J_i) > w_i/2$. By Lemma 4.4 we have $W(w_i,\vec{n}_i-\vec{n}'_{i-1}) \ge W(J_i) - 2\delta w_i \ge w_i/3$ for $\delta \le 1/12$. This finishes the proof.

(ii) Let $X_i = W(w_i,\vec{n}_i-\vec{n}'_{i-1})$ and let $Y_i = f(X_i/s_i)$ be the cost of the ith edge of the path. If $(w_{i-1},\vec{n}_{i-1}) = (w_i,\vec{n}_i)$, then $J_i=\emptyset$ and $f(W(J_i)) = f(0) = Y_i$. Otherwise by Lemma 4.4, $|W(J_i)-X_i| \le \delta(W(J_i)+2w_i) < \delta(1+\delta)(X_i+2w_i) < 8\delta X_i$; the last inequality follows since $X_i > w_i/3$ by the definition of an edge and $\delta < 1/12$. Thus $|C_i - X_i/s_i| = |W(J_i)/s_i - X_i/s_i| \le 8\delta X_i/s_i$ and by the condition (F$*$) and our choice of δ we get $|f(C_i)-Y_i| \le \varepsilon Y_i/3$. We get the required bound by summing, maximizing or minimizing over all edges of the path, as required by the type of the problem. ∎

We now finish the proof of Theorems 2.1 and 2.2 for the minimization versions; the case of maximization is similar. Let C^* be the optimal cost, let $C^\#$ be the cost of an optimal path in G, and let C be the cost of the output solution of the PTAS. By Lemma 3.1 there exists an optimal schedule with non-decreasing weights. Thus by Lemma 5.3 (i) its principal representation is a path in G, which cannot be cheaper than the optimal path. By Lemma 5.3 (ii) $C^\# \le (1+\varepsilon/3)C^*$. Using Lemma 5.3 (ii) for the output assignment we get

$$C \le \frac{1}{1-\frac{\varepsilon}{3}} C^\# \le \frac{1+\frac{\varepsilon}{3}}{1-\frac{\varepsilon}{3}} C^* \le (1+\varepsilon)C^*.$$

Thus we have found a required approximate solution.

6 Discussion

To get more insight in the meaning of the condition (F$*$), we prove the following characterization for convex functions.

Observation 6.1 *Suppose $f : [0, +\infty) \to [0, +\infty)$ is on $(0, +\infty)$ convex and not identically 0. Then f satisfies* (F$*$) *if and only if the following conditions hold:*

- *$f(x) > 0$ for any $x > 0$,*
- *for $x \to \infty$, $f(x)$ is polynomially bounded both from above and below (i.e., for some constant c, $f(x) \leq O(x^c)$ and $f(x) \geq \Omega(1/x^c)$).*
- *for $x \to 0$, $f(x)$ is polynomially bounded both from above and below (i.e., for some constant c, $f(x) \leq O(1/x^c)$ and $f(x) \geq \Omega(x^c)$).*

Proof: If $f(x) = 0$ for $x > 0$ then (F$*$) holds if and only if f is identically 0 on $(0, +\infty)$. Otherwise let $h_f(z) = \ln f(e^z)$ be as in (F$**$), note that it is defined everywhere. Since f is convex, it has both left and right derivatives everywhere (and they are equal almost everywhere). Thus also h_f has both derivatives everywhere, and (F$**$) (i.e., uniform continuity of h_f) is equivalent to the statement that the derivatives of h_f are bounded by some constants, both from above and below. This is equivalent to the fact that h_f is bounded by linear functions, from above and below, and for z approaching both $-\infty$ and $+\infty$. This in turn is equivalent with f being polynomially bounded, both from above and below, and for x approaching both 0 and $+\infty$. ∎

This characterization is related to Conjecture 4.1 of Alon et al. [1] which we now disprove. The conjecture says that for a convex function f, and for the problem of minimizing $\sum f(C_i)$ on identical machines the following three conditions are equivalent: (i) it has a PTAS, (ii) it has a polynomial approximation algorithm with a finite performance guarantee, (iii) the heuristic LPT, which orders the jobs according to non-increasing processing times and schedules them greedily on the least loaded machine, has a finite performance guarantee.

We know that if (F$*$) holds, there is a PTAS (for a computable f) [1]. Observation 6.1 implies that if (F$*$) does not hold, then LPT does not have a finite performance guarantee; the proof is similar to Observation 4.1 of [1] (which says that no such algorithm exists for an exponentially growing function, unless $P = NP$, by a reduction to KNAPSACK).

Now consider $f(x) = x^{t(x)}$ where t is some slowly growing unbounded function; $t(x) = \log\log\log\log x$ works. It is easy to verify that any such f is convex and does not satisfy (F$*$). However, it is possible to find a PTAS on identical machines using the integer programming approach of [1]. The function f does not satisfy (F$*$) on $[0, +\infty)$, but it satisfies (F$*$) for any interval $[0, T]$, moreover for a fixed ε the value of δ can be bounded by $\varepsilon / O(t(T))$. The PTAS algorithm now proceeds in the following way. It computes the bound on the completion time T as the sum of all processing times and chooses δ and

$\lambda = 1/\delta$ accordingly. Since M is at most singly exponential in the size of the instance, λ is proportional to a triple logarithm of the instance size. Now we use the integer programming approach from [1]. Resulting algorithm has time complexity doubly exponential in λ, thus the complexity is bounded by the size of the instance. Thus the algorithm is polynomial.

Let us conclude by a few remarks about the problem of minimizing $\max f(C_i)$. It is easy to approximate it for any increasing f satisfying (F$*$): just approximate the minimum makespan, and then apply f to that. Thus our extension to bimodal functions is not very strong. However, our techniques apply to a wider range of functions f. Suppose for example that the function f is increasing between 0 and 1, and then again between c and $+\infty$, with an arbitrary behavior between 1 and c. Then it is possible to prove a weaker version of Lemma 3.1, saying that for some (almost) optimal schedule, for any $i < j$, $T_i < \mu T_j$ (if $T_i, T_j > 0$). The constant μ depends only on the function f. This is sufficient for the approximation scheme, if we redefine the edges in G to be the ones with weight $\mu w'/3$ rather than $w'/3$, and choose the other constants appropriately smaller. We omit the details and precise statement, since this extension of our results does not seem to be particularly interesting.

7 Scheduling with rejection

In this section we study the problem from Example 4 in the introduction.

Theorem 7.1

(i) *Let f be a non-negative computable function satisfying the conditions* (F$*$) *and* (G$*$). *Then the problem of scheduling with rejection on uniformly related machine with the objection to minimize the sum of weights of rejected jobs plus* $\sum f(C_i)$ *possesses a polynomial approximation scheme.*

(ii) *Let f be a non-negative computable bimodal function satisfying the condition* (F$*$). *Then the problem of scheduling with rejection on uniformly related machines with the objection to minimize the sum of weights of rejected jobs plus* $\max f(C_i)$ *possesses a polynomial approximation scheme.*

(iii) *If the machines are identical, then the same is true (in both cases above) even if f is computable and satisfies only the condition* (F$*$).

Proof: The proof is a modification of our general PTAS. We give only a brief sketch.

(i) First consider the objective penalty plus $\sum f(C_i)$. We modify the graph G used in our PTAS in the following way. We add n auxiliary levels between any two levels of the original graph, as well as after the last level. Each level again has nodes corresponding to all principal configurations, the target node now is the node $\alpha(J)$ on the last auxiliary level. The edges entering the original nodes and their values are as before. The edges entering the auxiliary levels are as follows. There is an edge from a configuration $(w, \vec{n})$ to $(w', \vec{n}')$ iff the following holds: $w < w'$, and $n'' = \mathrm{scale}_{w \to w'}(\vec{n})$ satisfies $\vec{n}'' \leq \vec{n}'$ and $n''_i = n'_i$ for all $i \leq \lambda^2/2$. (The last condition says that $(w', \vec{n}'' - \vec{n}')$ represents only sets of jobs with all processing times greater than $w'/2$.) The value of the edge is the smallest total penalty of a set of jobs represented by $(w', \vec{n}'' - \vec{n}')$. Additionally, there are edges between identical configurations, with value 0.

The size of the graph increases by a factor of n. It is easy to construct the nodes and edges of the graph. Given an edge entering an auxiliary level, we know the number of jobs of each size that should be rejected, so we just add penalties of the appropriate number of jobs with the smallest penalties. Thus also the values of the edges can be computed efficiently.

A schedule is represented by a sequence of nodes, one on each level, such that on the ith original level we use the configuration $\alpha(A)$ for the following set of jobs A: let B be the set of jobs scheduled on the first i machines, let w be the weight of the principal configuration of B. Then A consist of B together with all jobs j rejected in the schedule such that $p_j \leq w$. On the previous auxiliary level we use the configuration $\alpha(A')$, where A' is A minus all jobs scheduled on M_i. On the last level we use the target node. The other nodes are arbitrary (we are again mainly interested in the representations that are paths in the graph).

Given a path in the graph, we can construct a schedule represented by it as follows. We follow the path from the source. Upon traversing an edge entering an auxiliary level, such that its endpoints have different configurations, we reject a subset of jobs represented by the difference of the configurations (cf. the definition of the edge), with the smallest total penalty among such sets. As described above, this can be computed efficiently. We are guaranteed that these jobs were not scheduled or rejected so far, as their processing time is greater than the weight of the configuration at the beginning of the edge. The sum of the values of the edges entering the intermediate nodes is exactly equal to the penalty of rejected jobs, and the contribution of the machines is bounded as before. Thus the cost of the path is close to the cost of the schedule.

Given the optimal schedule, we need to show that it is represented by a path. As before, we may assume that the non-zero weights of machines are non-decreasing. Given the schedule, the nodes on the original levels of the graph and the preceding intermediate levels are exactly determined, and connected by edges. Given any $A \subseteq B$, it is easy to verify that the node $\alpha(A)$ of an original level is connected by a path with the node $\alpha(B)$ of the nth intermediate level after it. The total value of the edges of this path is always at most the total penalty of the jobs in $B - A$. (Note that here we may have to use more than one non-trivial edge, as we may be rejecting jobs in different weight ranges; this is the reason why we use n auxiliary levels.) The value of the machines in the optimal schedule is bounded as before, and the penalties are lower bounded by the cost of the corresponding edges. Thus the optimal schedule is not much smaller than the cost of the path. This finishes the proof for the case of minimizing the penalty plus $\sum f(C_i)$.

(ii) The case of minimizing the penalty plus $\max f(C_i)$ is somewhat different. The obstacle is that the cost of a path in the graph should be sum of the costs of edges on certain levels plus the maximum of the costs of edges on the remaining levels; for such a problem we are not able to use the usual shortest path algorithm.

Let M be some bound on $\max f(C_i)$. We use a similar graph as above, with the following modification. We include an edge entering an original node only if its value would be at most M; we set its value to 0. Now the cost of the shortest path is an approximation of the minimal penalty among all schedules with $\max f(C_i) \leq M$. More precisely, similarly as in Lemma 5.3, if there is a schedule with $\max f(C_i) \leq (1 - \varepsilon/3)M$ and total penalty P, the shortest path has cost at most P; on the other hand, from a path with cost P we may construct a schedule with $\max f(C_i) \leq (1 + \varepsilon/3)M$ and penalty P.

Now we solve the optimization problem by using the procedure above polynomially many times. Let s_{min} and s_{max} be the smallest and the largest machine speeds, respectively. Let p_{min} be the minimal processing time of a job, and let T be the total processing time of all jobs. In any schedule, any non-zero completion time is between $b = p_{min}/s_{max}$ and $B = T/s_{min}$. Now we cycle through all values $x = (1 + \delta)^i b$, $i = 0, 1, \ldots$, such that $x \leq B$; the constant δ is chosen by the condition (F$*$), as in Section 5. In addition, we consider $x = 0$. The number of such x is polynomial in the size of the number B/b (i.e. in $\log(B/b)$), which is polynomial in the size of the instance. For each x, we compute $M = f(x)$, and find a corresponding schedule with the smallest penalty P by the procedure above. (As a technical detail, we have to round each x to a sufficient precision so that the length of x is polynomial in the size of the instance; this is possible to do so that the ratio between successive values of x never exceeds $1 + 2\delta$, and that is sufficient.) We chose the best of these schedules, and possibly the schedule rejecting all jobs. Since the relative change between any two successive non-zero value of x is at most 3δ, the relative change between the successive values of M is at most $\varepsilon/3$ (by our choice of δ using (F$*$)), and we cover all relevant values of M with sufficient density.

(iii) follows from the proof of (i) and (ii) by noting that no condition is needed to order the machines. ■

Acknowledgements.

We thank Yossi Azar and Gerhard Woeginger for useful discussions and introduction to this problem. We are grateful to Tamir Tassa and anonymous referees for pointing out errors in the conference version of the paper [7] and other useful comments.

References

[1] N. Alon, Y. Azar, G. J. Woeginger, and T. Yadid. Approximation schemes for scheduling on parallel machines. *J. of Scheduling*, 1:55–66, 1998.

[2] Y. Azar and L. Epstein. Approximation schemes for covering and scheduling on related machines. In *Proc. of the 1st Workshop on Approximation Algorithms for*

Combinatorial Optimization Problems (APPROX '98), Lecture Notes in Comput. Sci. 1444, pages 39–47. Springer-Verlag, 1998.

[3] Y. Bartal, S. Leonardi, A. Marchetti-Spaccamela, J. Sgall, and L. Stougie. Multiprocessor scheduling with rejection. *SIAM J. Disc. Math.*, 13(1):64–78, 2000.

[4] A. K. Chandra and C. K. Wong. Worst-case analysis of a placement algorithm related to storage allocation. *SIAM J. Comput.*, 4:249–263, 1975.

[5] R. A. Cody and E. G. Coffman. Record allocation for minimizing expected retrieval costs on drum-like storage devices. *J. Assoc. Comput. Mach.*, 23:103–115, 1976.

[6] J. Csirik, H. Kellerer, and G. J. Woeginger. The exact LPT-bound for maximizing the minimum completion time. *Oper. Res. Lett.*, 11:281–287, 1992.

[7] L. Epstein and J. Sgall. Approximation schemes for scheduling on uniformly related and identical parallel machines. In *Proc. of the 7th Ann. European Symp. on Algorithms, Lecture Notes in Comput. Sci. 1643*, pages 151–162. Springer-Verlag, 1999.

[8] D. K. Friesen and B. L. Deuermeyer. Analysis of greedy solutions for a replacement part sequencing problem. *Mathematics of Operations Research*, 6:74–87, 1981.

[9] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP–completeness*. Freeman, 1979.

[10] R. L. Graham. Bounds for certain multiprocessor anomalies. *Bell System Technical J.*, 45:1563–1581, Nov. 1966.

[11] R. L. Graham. Bounds on multiprocessor timing anomalies. *SIAM J. Appl. Math.*, 17(2):416–429, 1969.

[12] D. S. Hochbaum and D. B. Shmoys. Using dual approximation algorithms for scheduling problems: Theoretical and practical results. *J. Assoc. Comput. Mach.*, 34:144–162, 1987.

[13] D. S. Hochbaum and D. B. Shmoys. A polynomial approximation scheme for scheduling on uniform processors: Using the dual approximation approach. *SIAM J. Comput.*, 17:539–551, 1988.

[14] E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan, and D. B. Shmoys. Sequencing and scheduling: Algorithms and complexity. In S. C. Graves, A. H. G. Rinnooy Kan, and P. Zipkin, editors, *Handbooks in Operations Research and Management Science, Vol. 4: Logistics of Production and Inventory*, pages 445–552. North-Holland, 1993.

[15] S. Sahni. Algorithms for scheduling independent tasks. *J. Assoc. Comput. Mach.*, 23:116–127, 1976.

[16] G. J. Woeginger. A polynomial time approximation scheme for maximizing the minimum machine completion time. *Oper. Res. Lett.*, 20:149–154, 1997.