Load balancing of temporary tasks in the ℓ_p norm

Yossi Azar [a,1], Amir Epstein [a,2], Leah Epstein [b,3]

[a] *School of Computer Science, Tel Aviv University, Tel Aviv, Israel.*

[b] *School of Computer Science, The Interdisciplinary Center, Herzliya, Israel.*

Abstract

We consider the on-line load balancing problem where there are m identical machines (servers). Jobs arrive at arbitrary times, where each job has a weight and a duration. A job has to be assigned upon its arrival to exactly one of the machines. The duration of each job becomes known only upon its termination (this is called temporary tasks of unknown durations). Once a job has been assigned to a machine it cannot be reassigned to another machine. The goal is to minimize the maximum over time of the sum (over all machines) of the squares of the loads, instead of the traditional maximum load.

Minimizing the sum of the squares is equivalent to minimizing the load vector with respect to the ℓ_2 norm. We show that for the ℓ_2 norm the greedy algorithm performs within at most 1.493 of the optimum. We show (an asymptotic) lower bound of 1.33 on the competitive ratio of the greedy algorithm. We also show a lower bound of 1.20 on the competitive ratio of any algorithm.

We extend our techniques and analyze the competitive ratio of the greedy algorithm with respect to the ℓ_p norm. We show that the greedy algorithm performs within at most $2 - \Omega(1/p)$ of the optimum. We also show a lower bound of $2 - O(\ln p/p)$ on the competitive ratio of any on-line algorithm.

Email addresses: `azar@tau.ac.il` (Yossi Azar), `amirep@tau.ac.il` (Amir Epstein), `lea@idc.ac.il` (Leah Epstein).

[1] Research supported in part by the Israel Science Foundation and by the IST program of the EU.

[2] Research supported in part by the Israel Science Foundation and by the Deutsch Institute.

[3] Research supported in part by the Israel Science Foundation.

1 Introduction

We are given m parallel identical machines and a number of independent jobs (tasks) arriving at arbitrary times; each job has a weight and a duration. A job should be assigned upon its arrival to exactly one of the machines based only on the previous jobs and without any knowledge on the future jobs, thus increasing the load on this machine by its weight for the duration of the job. The duration of each job becomes known only upon its termination (this is called temporary tasks of unknown durations). The load of a machine is the sum of the weights of the jobs assigned to it. For any ℓ_p norm we define the cost of an assignment for an input sequence of jobs as the maximum over time of the ℓ_p norm of the load vector. Specifically, the ℓ_∞ norm is the makespan (or maximum load) and the ℓ_2 norm is the Euclidean norm, which is equivalent to the sum of the squares of the loads. The goal of an assignment algorithm is to assign all the jobs so as to minimize the cost.

Consider for example the case where the weight of a job corresponds to the frequency with which it accesses a disk. Each job may see a delay that is proportional to the load on the machine it is assigned to. Then the *average* delay is proportional to the sum of squares of the machines loads (namely the ℓ_2 norm of the corresponding machines load vector) whereas the *maximum* delay is proportional to the maximum load.

We measure the performance of an on-line algorithm by its competitive ratio. An on-line algorithm is c-competitive if for each input the cost of the assignment produced by the algorithm is at most c times larger than the cost of the optimal assignment.

We first summarize **our results**.

- For the ℓ_2 norm, we show that the greedy algorithm is at most 1.493-competitive for any number of machines. In fact, for $m = 2$ the greedy algorithm is optimal and its competitive ratio is 1.145 and for $m = 3$ we can improve the competitive ratio to 1.453.
- For the ℓ_2 norm, we show that there is no on-line algorithm that for all m is 1.202-competitive.
- For the ℓ_2 norm and for any given m, we show a lower bound of $2/\sqrt{3} - O(\frac{1}{m})$ on the competitive ratio of any on-line algorithm and a lower bound of $2/\sqrt{3}$ for m divisible by 3.
- For the ℓ_2 norm, we show (an asymptotic) lower bound of 1.338 on the competitive ratio of the greedy algorithm.
- For the general ℓ_p norm (for any $p > 1$), we show that the greedy algorithm is at most $2 - \Omega(1/p)$-competitive for any number of machines.
- For the general ℓ_p norm (for any $p > 1$), we show (an asymptotic) lower bound of $2 - O(\ln p/p)$ on the competitive ratio of any on-line algorithm.
- For the general ℓ_p norm (for any $p > 1$), we show that for $m = 2$ the greedy algorithm is an optimal on-line algorithm.

Temporary tasks, ℓ_∞ norm: For the problem of on-line load balancing of temporary tasks an algorithm with competitive ratio $2 - \frac{1}{m}$ was proved for permanent tasks (tasks that never depart) by Graham (12); nevertheless, Graham's analysis holds also for temporary tasks. The results in (4) show that his algorithm is optimal by constructing a lower bound of $2 - \frac{1}{m}$ on the competitive ratio of any on-line algorithm.

Permanent tasks, ℓ_∞ norm: This is the classic ancient problem of scheduling jobs on identical machines minimizing the makespan (or maximum load). Graham (12) showed that the greedy load balancing algorithm is $2 - \frac{1}{m}$-competitive in this case. The greedy algorithm is an optimal on-line algorithm only for $m \leq 3$ (9).

Bartal et al. (6) were the first to show an algorithm whose competitive ratio is strictly below $c < 2$ (for all m). More precisely, their algorithm achieves a competitive ratio of $2 - \frac{1}{70}$. Later, the algorithm was generalized by Karger, Phillips and Torng (16) to yield an upper bound of 1.945. Subsequently, Albers (1) designed 1.923-competitive algorithm. Fleischer and Wahl (10) improved this result to a ratio of 1.9201.

Bartal et al. (7) showed a lower bound of 1.8370 for the problem. This result was improved by Albers (1) to 1.852 and then by Gormley et al. (11) to 1.853. The best lower bound currently known is due to Rudin (15), who showed a lower bound of 1.88.

Permanent tasks, ℓ_p norm: Chandra and Wong (8) were the first to consider the problem of minimizing the the ℓ_2 norm of the machine loads. They showed that if the jobs arrive in non-increasing weight order then the greedy algorithm yields a schedule whose cost is within $\sqrt{25/24}$ of the optimal cost. This result was slightly improved by Leung and Wei (17). Chandra and Wong (8) also considered the general ℓ_p norm (for any $p > 1$) and showed that the greedy algorithm on the sorted items achieves a constant performance bound. The constant depends on p and grows to $\frac{3}{2}$ when p grows to ∞. The problem of on-line load balancing in the general ℓ_p norm (for any $p > 1$) for permanent tasks was considered in (3). The results in (3) show that for the ℓ_2 norm, the greedy algorithm performs within $2/\sqrt{3}$ of the optimum, and no on-line algorithm achieves a better competitive ratio. For the ℓ_2 norm (3) also provided an on-line algorithm with competitive ratio $2/\sqrt{3} - \delta$, for some fixed δ, for any sufficiently large number of machines. For the general ℓ_p norm the results show that the competitive ratio of the greedy algorithm is $2 - \Theta((\ln p)/p)$.

Off-line results: Azar et al. (5) presented a polynomial time approximation scheme for the problem of off-line load balancing of temporary tasks (in the ℓ_∞ norm) in the case where the number of machines is fixed. For the case

in which the number of machines is given as part of the input (i.e, not fixed) they showed that no polynomial algorithm can achieve a better approximation ratio than $\frac{3}{2}$ unless $P = NP$.

For the problem of off-line load balancing of permanent tasks (in the ℓ_∞ norm), there is a polynomial time approximation scheme for any fixed number of machines (13; 18) and also for arbitrary number of machines by Hochbaum and Shmoys (14).

Off-line scheduling and load balancing of permanent tasks with respect to the ℓ_p norm has been considered in (2). The off-line minimization problem is known to be NP-hard in the strong sense. Alon et al. (2) provided a polynomial approximation scheme for scheduling jobs with respect to the ℓ_p norm for any $p > 1$. An example in which the optimal assignment for the sum of the squares is different from the optimal assignment in the ℓ_∞ norm is also given in (2).

2 Definitions and preliminaries

In the **load balancing problem** we are given m identical machines (servers) and a finite sequence of events. We denote the input sequence by $\sigma = \sigma_1, \ldots, \sigma_r$. Each event σ_i is an arrival or departure of a job (task). We view σ as a sequence of times, the time σ_i is the moment after the $i^{\underline{th}}$ event happened. We denote the weight of a job j by w_j, its arrival time by a_j and its departure time (which is unknown until it departs) by d_j. An on-line algorithm has to assign a job upon its arrival without knowing the future jobs and the durations of jobs that have not departed yet. We compare the performance of on-line algorithms and an optimal off-line algorithm that knows the sequence of jobs and their durations in advance. Let $J_i = \{j \mid a_j \leq \sigma_i < d_j\}$ be the active jobs at time σ_i. A **schedule** S assigns each job j to a single machine k, $1 \leq k \leq m$. For every schedule S, the **load** of machine k at time σ_i, denoted $L_k^i(S)$, is the sum of weights of all jobs assigned to machine k in S, and active at this time. The **vector of loads at time** σ_i is $L^i(S) = (L_1^i(S), \ldots, L_m^i(S))$. Our cost measure is the ℓ_p norm. Hence the **cost** of a schedule S at time σ_i is defined as $\|L^i(S)\|_p = (\sum_{k=1}^m (L_k^i(S))^p)^{\frac{1}{p}}$. The **cost** of a schedule S is defined as $\|L(S)\|_p = \max_i \|L^i(S)\|_p$. We denote the load vector with the maximum cost, by $L(S) = (L_1(S), \ldots, L_m(S))$.

The **optimal cost**, denoted $OPT(S)$, is the minimal cost over all possible schedules for the given sequence of events.

We measure the performance of our algorithms by the competitive ratio. For a fixed $p > 1$, the competitive ratio of a schedule S is defined as $C(S) = \|L(S)\|_p / OPT(S)$. Let A be an on-line assignment algorithm. The competitive

ratio of A for a fixed number $m \geq 1$ of machines is defined as

$$C_{A,m} = \sup\{C(S) \mid S \text{ is a schedule produced by } A \text{ on } m \text{ machines}\}.$$

The competitive ratio of A for an arbitrary number of machines is defined as $C_A = \sup\{C_{A,m} \mid m \geq 1\}$.

The previous definitions cover also the case where we measure the sum of squares of loads, since then the cost is $(\|L(S)\|_2)^2$. Consequently, the competitive ratios for the sum of the squares of loads are equal to $(C(S))^2$, $(C_{A,m})^2$ and $(C_A)^2$ w.r.t. the ℓ_2 norm.

Now we define the notion of the shape of a schedule, which is an abstraction of a schedule where for every machine all jobs assigned to it except for one are replaced by very small identical jobs with the same total load. In general it may be the case that the schedule we analyze cannot be produced by the original algorithm which we study. Nevertheless, it upper bounds the cost of the original algorithm. Moreover, the concept of a shape is very useful for proving upper bounds on the competitive ratio, since the optimal assignment may improve (by partitioning the jobs) while the cost of the assignment does not change. Hence a shape is a pessimistic estimate of a schedule. A shape characterizes each machine by two numbers: a_i the total load of the small jobs, and u_i (a lower bound on) the weight of one large job. We denote a **shape** by a pair $R = (a, u)$, where a and u are vectors of m nonnegative reals. The **vector of loads** of a shape is defined as $L(R) = a + u$. The competitive ratio of a shape R is $C(R) = \|L(R)\|_p / OPT(R)$.

It is possible to compute the optimal cost of the shape $R = (a, u)$ explicitly. It is the cost of a schedule in which some big jobs are scheduled each on a separate machine and the rest of the jobs are balanced evenly on the rest of the machines. Let the machines be ordered so that u_i are nondecreasing. For $1 \leq l \leq m$ let $h_l = (\sum_{i=1}^{m} a_i + \sum_{i=1}^{l} u_i)/l$. Let k be the largest l such that $h_l \geq u_l$ (k is always defined, since $h_1 \geq u_1$). We define the **height** of the shape to be $h(R) = h_k$.

It is easy to see that a good candidate for an optimal schedule for the shape R is to put on each machine one job of size exactly u_i and partition a_i into a few jobs so that they can be balanced exactly on the k machines; then the load vector is $(h_k, \ldots, h_k, u_{k+1}, \ldots, u_m)$. See the Figures 1 and 2 for examples where $a_i = 1$ for all i.

Note that any shape of a schedule has the same cost as the cost of the schedule. In addition the optimal cost of the shape is not greater than the optimal cost of the schedule.

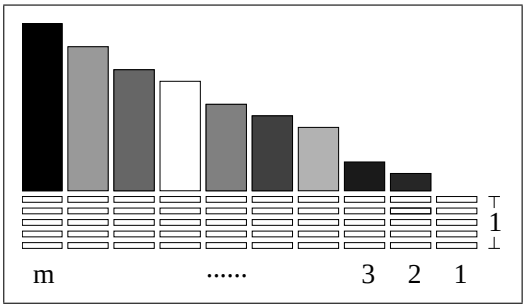

Fig. 1. A shape R.

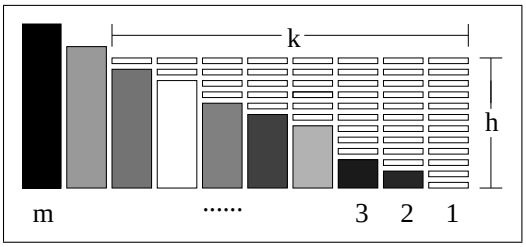

Fig. 2. Optimal assignment of R

Now we extend the notion of shape and define continuous shapes, which are defined similarly to shapes. This extension treats the machines as points in the interval $[0, m]$. The load vector is a function defined on that interval, and it is the sum of two functions. One function gives the total load of the small jobs on each machine. The other one gives the load of one big job on each machine. Formally, a **continuous shape** is a pair $R = (a, u)$, where a and u are functions (not necessarily continuous), defined in the interval $[0, m]$. The **function of loads** of a shape is defined as $L(R) = a+u$. Where $a(t)$ represents the total load of small jobs of equal size at point t in the interval $[0, m]$ and $u(t)$ represents the load of one big job at point t in the interval $[0, m]$. From the convexity of the function x^p it follows that the **optimal cost of a continuous shape** R is obtained by selecting some big jobs and assigning each one to a separate machine, and the remaining jobs are balanced evenly on the rest of the machines. Formally w.l.o.g let $u(t)$ be a non-decreasing function. There exist functions $u'(t)$ and $a'(t)$ such that $R' = (a', u')$ gives the optimal load $L(R')$ for the shape R. These functions are as follows: $u'(t) = u(t)$ and

$$a'(t) = \begin{cases} 0 & t \leq t' \\ \frac{1}{m-t'}(\int_0^m a(t)dt + \int_{t'}^m u(t)dt) - u(t) & t' < t \end{cases}$$

for some value t', s.t for $t_1 \leq t'$ and $t_2 > t'$, it holds that $u'(t_1) \geq a'(t_2) + u'(t_2)$.

The Transition from a shape to a continuous shape is defined as follows. Let $R = (a, u)$ be a shape, then its continuous shape $R' = (a', u')$ is

$$a'(t) = a_i \quad i - 1 < t \leq i\,,$$

$$u'(t) = u_i \quad i - 1 < t \leq i\,.$$

Note that in the above transition from a shape to a continuous shape the cost of the assignment does not change while the cost of the optimal assignment can only decrease.

3 The greedy algorithm

In this section we analyze the competitive ratio of the greedy algorithm defined below.

> **Algorithm Greedy**: Upon arrival of a job j assign it to the machine with the current minimum load (ties are broken arbitrarily).

Note that a job departs from a machine it was assigned to.

To obtain a constant upper bound for the performance ratio of Greedy, we show that each schedule can be replaced by a very special continuous shape so that the competitive ratio does not decrease. Computing the competitive ratio is then shown to be equivalent to computing the maximum of a certain function with equality constraints over the reals. A shape $R = (h, x)$ is called **partially flat** if there exists an integer $1 \leq k \leq m$, and a real $c > 0$ such that the following conditions hold:

$$\begin{aligned} h_i &= c && \text{for } i = 1, \ldots, k \\ x_i &\geq 0 && \text{for } i = 1, \ldots, k \\ h_i &= x_i = 0 && \text{for } i = k+1, \ldots, m. \end{aligned}$$

When $k = m$ the shape is called **flat**. A shape $R = (h, x)$ is called **separate** if there exists an integer $1 \leq k \leq m$ and a real $c > 0$, such that the following conditions hold:

$$\begin{aligned} h_i &= 0 \text{ for } i = 1, \ldots, k \\ x_i &\geq c \text{ for } i = 1, \ldots, k \\ h_i &= c \text{ for } i = k+1, \ldots, m \\ x_i &= 0 \text{ for } i = k+1, \ldots, m. \end{aligned}$$

Let S be a schedule obtained by Greedy and let $L(S)$ be the load when Greedy reaches the maximum cost. Let h be the load vector of all jobs except the last job assigned to each machine, we treat these jobs as very small jobs and call them **sand jobs**. Let x be the weight vector of the last job assigned to each machine. The shape $R = (h, x)$ is a shape of the schedule S, we call it the **Greedy** shape. Recall that when moving from a schedule to a shape and from a shape to a continuous shape the cost of the assignment does not change while the cost of the optimal assignment can only decrease.

Lemma 1 *Let $R = (h, x)$ be the Greedy shape of a Greedy schedule S, normalized by multiplying the weight of all the jobs by the same number such that $(OPT(S))^p = m$ [4]. Then there exists a partially flat shape $R' = (h', x')$ such that $\|L(R)\|_p \leq \|L(R')\|_p$ and $OPT(R) \geq OPT(R')$ with the non-zero components of h' equal to 1, and the off-line shape of the shape R' is a separate shape.*

Proof: It is easy to see that $h_i \leq 1$ (otherwise when the last job was assigned to machine i before Greedy reached the maximum cost, $(OPT(S))^p > m$, which is a contradiction).

W.l.o.g we assume that the machines are ordered so that $h_i + x_i$ are non increasing. We perform the following transformation. Note that in each step of the transformation the cost of Greedy can only increase and the cost of the off-line algorithm can only decrease. This is in particular due to the convexity of the function x^p. Figure 3 shows the shape resulting after each transformation step. In this figure the black area represents regular jobs and the white area represents sand jobs.

(1) In this step we move to a continuous shape $R = (h(t), x(t))$, where $h(t)$ and $x(t)$ are functions in the interval $[0, m]$. We treat each point t in that interval as a machine and the value $h(t) + x(t)$ as its load. Now we transform regular jobs (jobs or parts of jobs) that are placed below height 1 (where height represents machine load) into sand jobs. Next we push the sand jobs to the left such that the height of the sand jobs will be equal to 1 from point 0 to point V_0, where V_0 is the total volume of the sand jobs. See Step 1 in Figure 3 for the resulting shape of this step. Formally, let $R = (h(t), x(t))$ be the current shape and let t_0 be maximal value of t, such that $h(t)+x(t) \geq 1$ than the new shape $R' = (h'(t), x'(t))$ is obtained as follows. Denote

$$V_0 = t_0 + \int_{t_0}^{m} (h(t) + x(t))dt$$

then

$$h'(t) = \begin{cases} 1 & t \leq V_0 \\ 0 & V_0 < t \end{cases}$$

$$x'(t) = \begin{cases} h(t) + x(t) - 1 & t \leq t_0 \\ 0 & t_0 < t \end{cases}$$

(2) We would like to separate the machines of the off-line algorithm that have sand and possibly regular jobs into those having sand jobs only, and machines having regular jobs only. These are machines which all

[4] The number is $m^{1/p}$ divided by the original optimal cost of the schedule S.

of them have the same height. Consider regular jobs that are scheduled (by the off-line algorithm) on such machines together with sand jobs. Take the total area of these jobs and push it to the left. In this way, some machines will have only regular jobs of the same height as these machines had before this transformation. The machines on the right hand side will have the same height, but will contain only sand jobs, with the area of sand jobs just before this transformation. See Step 2 in Figure 3 for the resulting shape of this step. Formally, let $R = (h(t), x(t))$ be the current shape then the new shape $R' = (h'(t), x'(t))$ is obtained as follows. Let t_1 be a minimal point such that machine t_1 has sand jobs in the off-line algorithm, let w be its total load in the off-line algorithm and let V_1 be the volume of the regular jobs on machines $[t_1, m]$. Denote

$$V_1 = \int_{t_1}^{m} x(t)dt$$

then

$$h'(t) = h(t)$$

$$x'(t) = \begin{cases} x(t) & t \leq t_1 \\ w & t_1 < t \leq t_1 + \frac{V_1}{w} \\ 0 & \text{otherwise} \end{cases}$$

(3) Part of the sand jobs on machines with no regular jobs are transformed into regular jobs of height equal to w (as defined in the previous step). These jobs are assigned to the remaining machines that have only sand jobs. This is done such that all non empty machines in the Greedy shape have sand jobs and a regular job. See Step 3 in Figure 3 for the resulting shape of this step. Formally, let $R = (h(t), x(t))$ be the current shape then the new shape $R' = (h'(t), x'(t))$ is obtained as follows. Let t_1 be maximal such that machine t_1 has sand jobs and a regular job. Let t_2 be maximal such that machine t_2 has jobs (any jobs). Let $s = \frac{t_1 \cdot w + t_2}{w+1}$, then

$$h'(t) = \begin{cases} 1 & t \leq s \\ 0 & \text{otherwise} \end{cases}$$

$$x'(t) = \begin{cases} x(t) & t \leq t_1 \\ w & t_1 < t \leq s \\ 0 & \text{otherwise} \end{cases}$$

It is easy to see that in each of the transformation steps the cost of Greedy can only increase and the cost of the off-line algorithm can only decrease due

to the convexity of the function x^p. We denote the shape obtained by the transformation by $R' = (h', x')$. This shape has jobs on machines $[0, s]$ for some real number $0 < s < m$. Each machine $t \in [0, s]$ has sand jobs of height 1 and a regular job of height $x'(t) > 0$, other machines have no jobs assigned to them, hence this is a partially flat shape. The off-line shape has jobs of height $x'(t)$ on machines $t \in [0, s]$ and sand jobs of total volume s evenly assigned to machines $(s, m]$. In addition $\min_{t \in [0,s]} x'(t) \geq w = s/(m - s)$, hence the off-line shape is a separate shape.

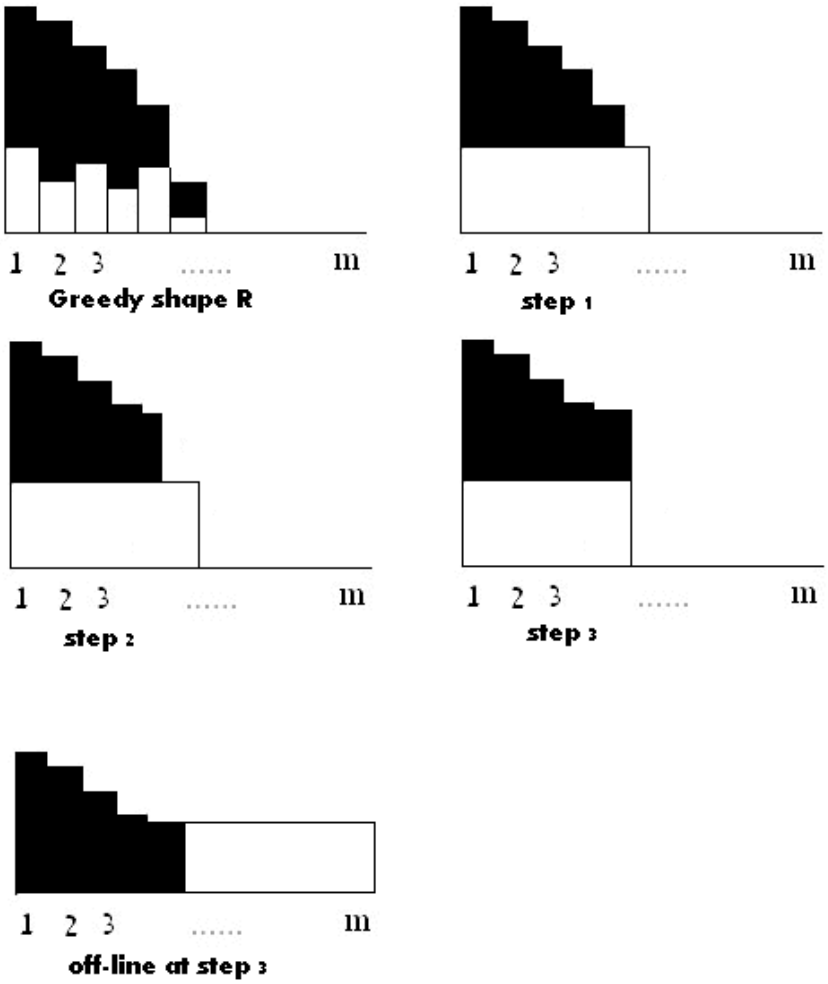

Fig. 3. The transformation steps

■

Lemma 2 *Let $R = (h, x)$ be a partially flat shape such that $h(t) = 1$ for $0 \leq t \leq s$. Assume that the off-line shape of R is a separate shape. Then there is a shape R'=(h,x'), such that for $0 \leq t \leq s$, $x'(t) = y$ for some value $y > 0$ and $x'(t) = 0$ for $t > s$, and it holds that $\|L(R)\|_p \leq \|L(R')\|_p$ and $OPT(R) = OPT(R')$.*

Proof: Let $\delta \cdot m = s$. Define

$$y = (\frac{1}{\delta \cdot m} \int_0^{\delta \cdot m} x^p(t)dt)^{\frac{1}{p}}.$$

Clearly $OPT(R) = OPT(R')$. Now

$$\begin{aligned}
\|L(R)\|_p &= (\int_0^{\delta \cdot m} (1 + x(t))^p dt)^{\frac{1}{p}} \\
&\leq (\int_0^{\delta \cdot m} 1dt)^{\frac{1}{p}} + (\int_0^{\delta \cdot m} x^p(t)dt)^{\frac{1}{p}} \\
&= (\delta \cdot m)^{\frac{1}{p}} + (\delta \cdot m)^{\frac{1}{p}} \cdot y \\
&= (\delta \cdot m)^{\frac{1}{p}} \cdot (1 + y) \\
&= \|L(R')\|_p
\end{aligned}$$

where the inequality follows from the triangle inequality for the ℓ_p norm. ■

Define the function $f(\delta, x)$ with the constraint $g(\delta, x)$.

$$f(\delta, x) = \delta \cdot (1 + x)^p, \tag{1}$$

$$g(\delta, x) = \delta \cdot x^p + \frac{\delta^p}{(1 - \delta)^{p-1}} = 1. \tag{2}$$

Theorem 3 *The competitive ratio of the greedy algorithm satisfies*

$$C_{Greedy} \leq (f_{max})^{\frac{1}{p}}$$

where f_{max} is the maximum of f with constraint (2), in the domain $0 \leq x, 0 \leq \delta \leq \frac{1}{2}$.

Proof: Let R_0 be the Greedy shape of a schedule S obtained by Greedy. For simplicity we transform it to a new shape R_1 by normalizing all job weights such that

$$(OPT(R_1))^p = m. \tag{3}$$

If all the h components of R_1 are equal to zero then the Greedy schedule is optimal. Otherwise by Lemma 4 and Lemma 2 we obtain from R_1 a partially flat shape $R_2 = (h, x)$, in which all the non zero components of x are the same and its off-line shape is a separate shape such that $\|L(R_1)\|_p \leq \|L(R_2)\|_p$ and $OPT(R_1) \geq OPT(R_2)$. We have

$$(\|L(R_2)\|_p)^p = \delta \cdot m \cdot (1+x)^p, \tag{4}$$

$$(OPT(R_2))^p = \int_0^{\delta \cdot m} x^p(t)dt + (1-\delta)m(\frac{\delta \cdot m}{(1-\delta)\cdot m})^p = \delta \cdot m \cdot x^p + \frac{\delta^p \cdot m}{(1-\delta)^{p-1}}, \tag{5}$$

$$x \geq \frac{\delta \cdot m}{(1-\delta)\cdot m} = \frac{\delta}{1-\delta}. \tag{6}$$

The last inequality restricts the weight of a regular job to be greater than the total weight of sand jobs on machines with sand jobs in the off-line shape. For simplicity we divide equalities (4) and (5) by m, this does not change the ratio between the cost and the off-line cost of shape R_2, which gives the following

$$f(\delta, x) = \delta \cdot (1+x)^p, \tag{7}$$

$$1 \geq \delta \cdot x^p + \frac{\delta^p}{(1-\delta)^{p-1}}, \tag{8}$$

$$x \geq \frac{\delta}{1-\delta}. \tag{9}$$

The first inequality results from (5), since $(OPT(R_2))^p \leq (OPT(R_1))^p = m$ and the division by m.

Substituting (9) in (8) gives

$$1 \geq \frac{\delta^{p+1}}{(1-\delta)^p} + \frac{\delta^p}{(1-\delta)^{p-1}} = \frac{\delta^p}{(1-\delta)^p}$$

which yields $\delta \leq \frac{1}{2}$. We obtain the following relation for f

$$f(\delta, x) = \frac{\frac{1}{m}(\|L(R_2)\|_p)^p}{\frac{1}{m}\cdot m} \geq \frac{(\|L(R_1)\|_p)^p}{(OPT(R_1))^p} \geq \frac{(\|L(S)\|_p)^p}{(OPT(S))^p}$$

where the first inequality follows from (3) and the fact that $\|L(R_1)\|_p \leq \|L(R_2)\|_p$. Hence to bound the competitive ratio of Greedy we need to solve the following maximization problem in the domain $0 \leq \delta \leq \frac{1}{2}$ and $0 \leq x$. We need to find the maximum of $f(\delta, x)$ under the constraint (8). It is easy to see that the maximum of f is obtained when (8) is an equality. ■

The next theorem follows from Theorem 3.

Theorem 4 *For the ℓ_2 norm the competitive ratio of the greedy algorithm is $C_{Greedy} \leq 1.493$.*

Proof: By Theorem 3 the competitive ratio of Greedy is obtained by solving the following maximization problem for f, which is obtained by substituting $p = 2$ in (1) and (2).

$$\text{maximize} \quad f(\delta, x) = \delta \cdot (1+x)^2, \tag{10}$$

$$\text{subject to} \quad 1 = \delta \cdot x^2 + \frac{\delta^2}{1-\delta}. \tag{11}$$

We solve this maximization problem. (11) gives

$$x = \sqrt{\frac{1 - \frac{\delta^2}{1-\delta}}{\delta}}. \tag{12}$$

Substituting equation (12) in equation (10) gives

$$f(\delta) = \delta \cdot \left(1 + \sqrt{\frac{1 - \frac{\delta^2}{1-\delta}}{\delta}}\right)^2. \tag{13}$$

Using Maple we found that the maximum of f is achieved at $\delta \approx 0.3642$, ($x \approx 1.474$), and $(C_{Greedy})^2 \leq f(\delta) \approx 2.2293$. ■

We note that for ℓ_2 norm and 3 machines the upper bound can be improved to 1.453 using a more detailed computation.

Now we turn to the case of the general ℓ_p norm.

Theorem 5 *For any $p > 1$, $C_{Greedy} \leq 2 - \Omega\left(\frac{1}{p}\right)$.*

Proof: By Theorem 3 we have

$$1 = g(\delta, x) \geq \delta \cdot x^p.$$

Hence

$$\delta \leq \frac{1}{x^p}.$$

Substituting δ in (1) gives

$$f(\delta, x) \leq \frac{(1+x)^p}{x^p} = f_1(x). \tag{14}$$

Substituting $\delta \leq \frac{1}{2}$ in (1) gives

$$f(\delta, x) \leq \frac{1}{2}(1+x)^p = f_2(x).$$

Since $f_1(x)$ is a monotonically decreasing function of x and $f_2(x)$ is a monotonically increasing function of x we obtain

$$f(\delta, x) \leq f_{max} \leq f_2(x_0 = 2^{\frac{1}{p}}) = \frac{1}{2}\left(1 + 2^{\frac{1}{p}}\right)^p,$$

where $x_0 = 2^{\frac{1}{p}}$ is a solution of the equation

$$f_1(x) = f_2(x).$$

Hence

$$C_{Greedy} \leq (f_{max})^{\frac{1}{p}} \leq \frac{1+2^{\frac{1}{p}}}{2^{\frac{1}{p}}} = 1 + \left(\frac{1}{2}\right)^{\frac{1}{p}} = 1 + e^{-\frac{1}{p}\ln 2} = 2 - \Omega\left(\frac{1}{p}\right).$$

■

Next we consider the case of 2 machines. We use the following notation

$$D(p) = \sup_{x \geq 0} \left(\frac{1+(1+x)^p}{2^p + x^p}\right)^{\frac{1}{p}}. \tag{15}$$

It is proved in (3) that the supremum is achieved as a maximum at the unique solution $x \in (0, \infty)$ of the equation

$$x^{p-1}(1 + (1+x)^{1-p}) = 2^p.$$

Theorem 6 *For $m = 2$ and for any $p > 1$ the greedy algorithm is optimal and its competitive ratio is*

$$C_{Greedy,2} = D(p) \tag{16}$$

Proof: The lower bound follows immediately from the case of permanent tasks (see (3)). Next we prove the upper bound. The proof requires the following Lemma.

Lemma 7 *For $m = 2$ let S be a schedule obtained by Greedy. Then there exists a flat shape $R = (a, u)$ which is a shape of S.*

Proof: Consider time T, when Greedy reaches the maximum cost. Let $L = L(S)$ be the vector of loads of S at time T. W.l.o.g we assume that L_1 is the smallest component of L. We claim that the shape $R = (a, u)$ where $a = L_1$ and $u_i = L_i - a$, is a flat shape of $L(S)$. Clearly $L(R) = L(S)$. Consider machine 2 with $u_2 > 0$. Let j be the last job assigned to the machine 2 until time T. At the time of its assignment, the load of machine 1 must have been at most a, as otherwise the Greedy cost at that time would be greater than $\|L(S)\|_p$, which is a contradiction. Hence $w_j \geq L_2 - a = u_2$, and the shape $R = (a, u)$ is a flat shape. ■

We also need the following Lemma which is implicit in (3).

Lemma 8 *Let $R = (a, u)$ be a flat shape for $m = 2$. Then for any $p > 1$, $C(R) = D(p)$.*

To complete the proof of Theorem 6 we first apply Lemma 7 to obtain a flat shape R of S, which satisfies $C(S) \leq C(R)$. Next we apply Lemma 8 to obtain $C(R) = D(p)$.

■

From the above theorem it follows that for $m = 2$ and for the ℓ_2 norm Greedy is optimal and its competitive ratio is $D(2) = \sqrt{(\sqrt{5}+3)/4} \approx 1.145$, where the value of $D(2)$ is obtained from (3).

4 Lower bounds

4.1 Lower bounds for ℓ_2 norm

In this section we give lower bounds for the ℓ_2 norm. We prove a lower bound for any algorithm (the proof is for $m = 3$). Then we prove a weaker lower bound for any $m \geq 3$. Finally we prove a lower bound for Greedy for a large number of machines ($m \to \infty$).

Theorem 9 *For any on-line assignment algorithm A, it holds that $C_A \geq C_{A,3} \geq 1.202$.*

Proof: Consider the following sequence for three machines. First three unit jobs and one job of weight $x \geq 1$ arrive. Then two unit jobs depart. At last one job of weight 2 arrive. Consider the first three unit jobs. If algorithm A assigns two or more jobs to the same machine, it does not get any other job. Its cost is at least 5, the optimal cost is 3, and we are done. Otherwise, algorithm A assigns one unit job to every machine (the off-line algorithm assigns two unit jobs to machine 1 and one unit job to machine 2). Now the next job of weight x arrives. Algorithm A assigns it to one of the machines say 1 (the off-line algorithm assigns it to machine 3). Then two unit jobs depart, which are the jobs on machines 2,3 (the jobs on machine 1 in the off-line algorithm). At last a job of weight 2 arrive. The best algorithm A can do is to assign it to one of the empty machines 2 or 3 (the off-line algorithm assigns it to machine 1). Its cost is at least $\sqrt{(1+x)^2 + 2^2}$, whereas the optimum cost is $\sqrt{2^2 + 1 + x^2}$. The maximal ratio ≈ 1.202 is achieved for $x = \sqrt{5}$. ■

Theorem 10 *For any number of machines $m \geq 3$ and any on-line assignment algorithm A, it holds that $C_{A,m} \geq 2/\sqrt{3} - O(\frac{1}{m})$. For m divisible by 3, it holds that $C_{A,m} \geq 2/\sqrt{3}$.*

Proof: Let $m = 3k$. We consider the following sequence. First $4k$ unit jobs arrive. Then $2k$ unit jobs depart. Finally k jobs of weight 2 arrive. Consider the arrival of the first $4k$ unit jobs. Algorithm A assigns these jobs. W.l.o.g we assume that machines $1, \ldots, m$ are sorted in nondecreasing order of load (the off-line algorithm assigns two jobs to each machine $1, \ldots, k$ and one job to each machine $k+1, \ldots, 3k$). Then $2k$ jobs depart. There exists a minimal $t \geq 2k$, such that machines $1, \ldots, t$ are assigned at least $2k$ jobs. Then $2k$ jobs from machines $1, \ldots, t$ depart as follows, all jobs from machines $1, \ldots, t-1$ and some jobs from machine t (in the off-line algorithm the jobs on machines $1, \ldots, k$ depart). At the end of this step machines $1, \ldots, 2k$ are empty. Next k jobs of weight 2 arrive. The best algorithm A can do is to assign each job to an empty machine (In the off-line algorithm these jobs are assigned to machines $1, \ldots, k$). Finally there are jobs of total weight $4k$ assigned to no more than $2k$ machines. Due to the convexity of the function x^p, the minimum cost is obtained when all machines have the same load, therefore its cost is at least $\sqrt{2k \cdot (2)^2}$. The optimum cost is $\sqrt{k \cdot 2^2 + 2k \cdot 1^2}$, which yields a ratio of $2/\sqrt{3}$. For m not divisible by 3 a similar proof gives a ratio of $2/\sqrt{3} - O(\frac{1}{m})$.

■

Theorem 11 *For the greedy algorithm, $C_{Greedy} \geq 1.338$.*

Proof: First we prove a weaker lower bound of 1.314 for the greedy algorithm, by solving an ordinary differential equation analytically. Then by a similar proof we obtain a more complex ordinary differential equation, which has no simple solution. Hence we use a computer program to compute the competitive ratio in this case, which gives a lower bound of 1.338.

We start with the first proof. We see the m machines as m points in the interval $(0, 1]$, machine i as the point $\frac{i}{m} \in (0, 1]$, and the load of the machines as a function $f(t)$, $f(t) = l_i$ for $\frac{(i-1)}{m} < t \leq \frac{i}{m}$, where l_i is the load of machine i. For each machine i the total load is the value of f in the interval $(\frac{i-1}{m}, \frac{i}{m}]$ and the total load of all machines is the total volume of f in the interval $(0, 1]$ multiplied by m. Let $f(k/m)$ be the load of machine k at the end of step k and let $F(k/m)$ be the volume of the jobs assigned to machines $k, \ldots, m$ at the beginning of step k in the following process. For convenience we number the steps from m to 1 in decreasing order. In this process we keep the volume of jobs fixed and equal to 1 at the end of each step. We start with the arrival of infinitesimally small jobs of total volume 1, we call jobs of this type sand jobs. Both the off-line and the greedy algorithms assign these jobs evenly on all the machines (total height 1 on each machine). At step k a job of height x ($x \geq 1$)

arrives. Greedy assigns this job to the machine with minimum load w.l.o.g to machine k which is the one with the largest index among all machines with the minimum load $1, \ldots, k$ (otherwise we swap indices) and the off-line algorithm performs the same assignment. Then the sand jobs on machines $1, \ldots, k-1$ depart in Greedy. In the off-line algorithm the departing jobs are composed of all the sand jobs of machine k and equal amounts of sand jobs from machines $1, \ldots, k-1$ with the appropriate volume. Next sand jobs arrive with total volume $1 - F(k/m) - \frac{x}{m}$ ($= 1 - \textit{total}$ volume of machines $k, \ldots, m$), thus keeping the total volume equal to 1. Greedy and the off-line algorithm assign the sand jobs to machines $1, \ldots, k-1$ evenly, such that these machines have the same load. At the end of step k

$$f(k/m) = \frac{1 - F(k/m)}{k/m} + x.$$

When $m \to \infty$ $t = k/m$ is a continuous variable in the interval $[0, 1]$ and we get the following equation

$$f(t) = \frac{1 - F(t)}{t} + x. \tag{17}$$

We have $f(t) = -\frac{dF(t)}{dt}$ (since $F(t) = \int_t^1 f(u)du$) and we get

$$-\frac{dF(t)}{dt} = \frac{1 - F(t)}{t} + x.$$

Now we have the following first order differential equation

$$\begin{aligned} -t \cdot \frac{dF(t)}{dt} + F(t) - x \cdot t - 1 &= 0 \\ F(1) &= 0 \end{aligned}$$

It is easy to verify its solution

$$F(t) = -x \cdot t \cdot \ln(t) - t + 1. \tag{18}$$

Substituting equation (18) in equation (17) gives

$$f(t) = x \cdot \ln(t) + x + 1. \tag{19}$$

The above process continues until assigning the job with weight x to the machine represented by t_0, where $F(t_0) = 1$, i.e until the volume of all machines of Greedy that were assigned a job of weight x approaches 1 (there are no sand jobs that can depart from machines $0 \leq t < t_0$). From (18) we get

$$-x \cdot t_0 \cdot \ln(t_0) - t_0 + 1 = 1,$$

which gives

$$t_0 = e^{-\frac{1}{x}}.$$

At the end of the above process each machine in the interval $[t_0, 1]$ has sand jobs and one big job of weight x. In the off-line algorithm each machine in the interval $[t_0, 1]$ has one job of weight x and the other machines have sand jobs equally distributed among them. The maximum cost of Greedy and the off-line algorithm is obtained at the end of the above process due to the convexity of the function x^p. Let $Greedy(x)$ and $Opt(x)$, be the costs of greedy and the off-line algorithms as functions of x respectively.

$$\begin{aligned}
Greedy^2(x) &= \int_{e^{-\frac{1}{x}}}^{1} f^2(t)dt \\
&= \int_{e^{-\frac{1}{x}})}^{1} (x \cdot \ln(t) + x + 1)^2 dt \\
&= \int_{e^{-\frac{1}{x}}}^{1} [x^2 \cdot \ln^2(t) + 2x \cdot (x+1) \cdot \ln(t) + (x+1)^2] dt \\
&= \Big[x^2 \cdot (t \cdot \ln^2(t) - 2t \cdot \ln(t) + 2t) + 2x \cdot (x+1) \cdot (t \cdot (\ln(t) - t) + (x+1)^2 \cdot t\Big]_{e^{-\frac{1}{x}}}^{1} \\
&= \Big[t \cdot (x^2 \cdot \ln^2(t) + 2x \cdot \ln(t) + x^2 + 1)\Big]_{e^{-\frac{1}{x}}}^{1} \\
&= x^2 \cdot (1 - e^{-\frac{1}{x}}) + 1,
\end{aligned}$$

$$\begin{aligned}
Opt^2(x) &= \left(1 - e^{-\frac{1}{x}}\right) \cdot x^2 + e^{-\frac{1}{x}} \cdot \left[\frac{1 - (1 - e^{-\frac{1}{x}}) \cdot x}{e^{-\frac{1}{x}}}\right]^2 \\
&= \frac{e^{-\frac{1}{x}} \cdot (1 - e^{-\frac{1}{x}}) \cdot x^2 + [1 - (1 - e^{-\frac{1}{x}}) \cdot x]^2}{e^{-\frac{1}{x}}} \\
&= e^{\frac{1}{x}} \cdot \left[(x^2 - 2x) \cdot (1 - e^{-\frac{1}{x}}) + 1\right].
\end{aligned}$$

Let

$$C(x) = \frac{Greedy(x)}{Opt(x)}$$

and

$$C = \max_{1 \le x} C(x).$$

For $x \ge 1$ the maximum value of $C(x)$ is obtained approximately at $x \approx 1.2612$ and its value is $C \approx 1.314$. Hence $C_{Greedy} \ge C \approx 1.314$.

Now we give a similar proof to improve the lower bound. The process is similar to the one described above with the difference that here we keep the cost of the off-line algorithm fixed and equal to 1 instead of keeping the volume fixed and equal to 1 at the end of each step. In this proof we use the same notations as in the first proof. Consider the off-line algorithm at the end of

step t when assigning a job with weight x to machine t. According to the invariant constraint we have.

$$(1-t)\cdot x^2 + t\cdot h^2 = Opt^2(x) = 1,$$

where h is the weight of the sand jobs on machines $[0,t]$ at the end of step t. We define h as a function of t. The above equation gives

$$h(t) = \sqrt{\frac{1-(1-t)\cdot x^2}{t}}.$$

At the end of step t

$$f(t) = \frac{(1-t)\cdot x + t\cdot h(t) - F(t)}{t} + x. \tag{20}$$

We have $f(t) = -\frac{dF(t)}{dt}$ and we get

$$-\frac{dF(t)}{dt} = \frac{(1-t)\cdot x + t\cdot h(t) - F(t)}{t} + x.$$

Now we have the following first order differential equation

$$\begin{aligned} -t\cdot\frac{dF(t)}{dt} + F(t) - t\cdot h(t) - x &= 0 \\ F(1) &= 0. \end{aligned}$$

The solution to this equation is not simple and was calculated using a computer program, which gave the following result. For $x \geq 1$ the maximum value of $C(x)$ is obtained approximately at $x \approx 1.3888$ and its value is $C \approx 1.338$. Hence $C_{Greedy} \geq C \approx 1.338$, which completes the proof. ■

4.2 Lower bound for general $p > 1$

In this section we construct a lower bound for general $p > 1$.

Theorem 12 *For any $p > 1$ and any on-line algorithm A, it holds that $C_A \geq 2 - O\left(\frac{\ln p}{p}\right)$.*

Proof: Let $m \to \infty$. As in the proof of Theorem 11 we consider the machines as points in the interval $(0,1]$, each machine is represented by a point $t \in (0,1]$, and the load of machine is represented as a function $f(t)$ in that interval. Let $0 < \alpha < 1$. We consider the following sequence. First sand jobs of total volume 1 arrive. Next, jobs of total volume $(1-\alpha)$ depart. Finally unit jobs of total volume 1 arrive. Consider the arrival of the sand jobs. Algorithm A

assigns these jobs, w.l.o.g we assume that machines $(0,\dots,1]$ are sorted non-increasingly by load (the off-line algorithm assigns these jobs evenly on all the machines). Let $s \le \alpha$ be maximal such that machines $(s,\dots,1]$ are assigned jobs of total volume $(1-\alpha)$. Then jobs of total volume $(1-\alpha)$ depart from machines $(s,\dots,1]$ (in the off-line algorithm these jobs depart evenly from all the machines). We denote by $x \le \alpha$ the fraction of machines with assigned jobs of total height greater than 1. Next the unit jobs of total volume 1 arrive. The best policy that Greedy can have at this point is to assign jobs of total volume $(1-\alpha)$ evenly to machines $(\alpha,\dots,1]$ and then to assign jobs of total volume α to the α least loaded machines. The least loaded machines are composed of machines that have jobs of total height less than 1 (i.e. machines $(x,\dots,\alpha]$) and machines that have jobs of total height 1 (w.l.o.g machines $(\alpha,\alpha+x]$). The off-line algorithm assigns these jobs evenly to all the machines. Let A and Opt be the costs of algorithm A and the off-line algorithms respectively.

$$\begin{aligned} A^p &\ge x\cdot 2^p + \alpha\left[\frac{\alpha+(\alpha-x)}{\alpha}\right]^p \\ &= x\cdot 2^p + \alpha\left(2-\frac{x}{\alpha}\right)^p, \end{aligned}$$

where the first term on the right hand side represents the cost of machines $(\alpha,\dots,\alpha+x]$ and the other term is a lower bound for the cost of machines $(0,\dots,\alpha]$.

$$Opt^p = (1+\alpha)^p.$$

Hence

$$C_A^p \ge \left(\frac{A}{Opt}\right)^p \tag{21}$$

$$\ge \frac{x\cdot 2^p + \alpha\left(2-\frac{x}{\alpha}\right)^p}{(1+\alpha)^p}. \tag{22}$$

We choose $\alpha = \frac{1}{p}$. We consider two cases. In both cases we use the inequality $e^{-x} \ge 1-x$. For $x \ge \frac{\alpha}{p} = \frac{1}{p^2}$ we obtain

$$C_A^p \ge \frac{\alpha}{p}\cdot\frac{2^p}{(1+\alpha)^p} = \frac{1}{p^2}\cdot\frac{2^p}{(1+\frac{1}{p})^p} = 2^p\frac{e^{-2\ln p}}{(1+\frac{1}{p})^p}.$$

Hence

$$C_A \ge 2\frac{e^{-2\frac{\ln p}{p}}}{1+\frac{1}{p}} \ge 2\frac{1-2\frac{\ln p}{p}}{1+\frac{1}{p}} = 2 - O\left(\frac{\ln p}{p}\right).$$

For $x < \frac{\alpha}{p} = \frac{1}{p^2}$ we obtain

$$C_A^p \geq \alpha \cdot \frac{(2-\frac{1}{p})^p}{(1+\alpha)^p} = \frac{1}{p} \cdot \frac{(2-\frac{1}{p})^p}{(1+\frac{1}{p})^p} = e^{-\ln p} \cdot \frac{(2-\frac{1}{p})^p}{(1+\frac{1}{p})^p}.$$

Hence

$$C_A \geq e^{-\frac{\ln p}{p}} \cdot \frac{2-\frac{1}{p}}{1+\frac{1}{p}} \geq (1-\frac{\ln p}{p}) \cdot \frac{2-\frac{1}{p}}{1+\frac{1}{p}} = 2 - O\left(\frac{\ln p}{p}\right).$$

■

References

[1] S. Albers. Better bounds for online scheduling. *SIAM Journal on Computing*, 29:459–473, 1999.

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

[3] A. Avidor, Y. Azar, and J. Sgall. Ancient and new algorithms for load balancing in the ℓ_p norm. *Algorithmica*, 29:422–441, 2001.

[4] Y. Azar and L. Epstein. On-line load balancing of temporary tasks on identical machines. *Siam Journal on Discrete Mathematics*, 18(2):347–352, 2004.

[5] Y. Azar, O. Regev, J. Sgall, and G. Woeginger. Off-line temporary tasks assignment. *Theoretical Computer Science*, 287:419–428, 2002.

[6] Y. Bartal, A. Fiat, H. Karloff, and R. Vohra. New algorithms for an ancient scheduling problem. *Journal of Computer and System Sciences*, 51(3):359–366, 1995.

[7] Y. Bartal, H. Karloff, and Y. Rabani. A better lower bound for on-line scheduling. *Information Processing Letters*, 50:113–116, 1994.

[8] A.K. Chandra and C.K. Wong. Worst-case analysis of a placement algorithm related to storage allocation. *SIAM Journal on Computing*, 4(3):249–263, 1975.

[9] U. Faigle, W. Kern, and G. Turan. On the performance of online algorithms for partition problems. *Acta Cybernetica*, 9:107–119, 1989.

[10] R. Fleischer and M. Wahl. Online scheduling revisited. *Journal of Scheduling*, 3(5):343–353, 2000.

[11] T. Gormley, N. Reingold, E. Torng, and J. Westbrook. Generating adversaries for request-answer games. In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 564–565, 2000.

[12] R. L. Graham. Bounds for certain multiprocessing anomalies. *Bell System Technical Journal*, 45:1563–1581, 1966.

[13] R.L. Graham. Bounds on multiprocessing timing anomalies. *SIAM J. Appl. Math*, 17:416–429, 1969.

[14] D. S. Hochbaum and D. B. Shmoys. Using dual approximation algorithms for scheduling problems: Theoretical and practical results. *J. Assoc. Comput. Mach.*, 34(1):144–162, January 1987.

[15] J.F. Rudin III. *Improved bounds for the on-line scheduling problem.* PhD thesis, The University of Texas at Dallas, May 2001.

[16] D. Karger, S. Phillips, and E. Torng. A better algorithm for an ancient scheduling problem. *Journal of Algorithms*, 20(2):400–430, 1996.

[17] J.Y.T. Leung and W.D. Wei. Tighter bounds on a heuristic for a partition problem. *Information Processing Letters*, 56:51–57, 1995.

[18] S. Sahni. Algorithms for scheduling independent tasks. *Journal of the Association for Computing Machinery*, 23:116–127, 1976.