Fairness-Free Periodic Scheduling

Jiri Sgall [*]   Hadas Shachnai [†]   Tami Tamir [‡]

Abstract

We consider a problem of *repeatedly* scheduling n jobs on m parallel machines. Each job is associated with a profit, gained each time the job is completed, and the goal is to maximize the average profit per time unit. Unlike other periodic scheduling problems, there is no fairness requirement. Still, it is impossible to process only the most profitable jobs, because once the processing of a job is completed, it goes on *vacation* and returns to the system, ready to be processed again, only after its vacation is over. This problem of *scheduling with vacations* has many applications, in production planning, machine maintenance, media-on-demand and databases query processing, among others.

We show that the problem is NP-hard already for jobs with unit processing times and unit profits, and develop approximation algorithms, as well as optimal algorithms for certain subclasses of instances. In particular, we show that a preemptive greedy algorithm achieves a ratio of 2 to the optimal for instances with arbitrary processing times and arbitrary profits. For the special case of unit processing times, we present a 1.66-approximation algorithm for instances with arbitrary profits, and a 1.39-approximation algorithm for instances with the same (unit) profits. For the latter case, we also show that when the load generated by an instance is sufficiently large (in terms of n and m), any algorithm that uses no intended idle times yields an optimal schedule.

1 Introduction

We consider a scheduling problem in which jobs need to be scheduled repeatedly. The input is a set of m identical parallel machines and a set of n jobs, $J = \{1, \ldots, n\}$, that need to be scheduled on the machines. All the jobs are ready at time $t = 0$. Each job, j, is associated with a processing time $p_j \geq 1$, and a window $a_j \geq p_j$; once the processing of j is completed, it goes on *vacation* and returns after $a_j - p_j$ time units, ready to be processed again. Thus, a feasible schedule is one where each machine processes at most one job at any time, and there is a gap of at least $a_j - p_j$ time units between two consecutive executions of job j. The jobs are sequential, i.e.,no job can be scheduled on more than one machine at the same time. There is also a profit, b_j, associated with each execution of job j; the goal is to find a feasible schedule

that maximizes the average profit per time unit. We call this problem *scheduling with vacations (SWV)*. Our problem arises in many practical scenarios where tasks need to be accomplished periodically, but due to setup times or maintenance requirements there must be some gap between two consecutive executions of a task. The following are two of the applications that motivate our study.

Surveillance Camera Scheduling: Consider a system of robots carrying surveillance cameras, which patrol an area periodically. Each robot has a predefined path that it needs to patrol, while recording all the events along this path. Upon completion of each patrol, the robot returns to the controller, where the recorded data is downloaded/processed, and the robot is prepared for its next tour. Each patrol j is associated with a profit b_j, gained once the corresponding robot completes the tour. The controller can handle at most m robots simultaneously, and it takes p_j time units to complete the processing of robot j. The time it takes robot j to traverse its path is $\ell_j \geq 1$. The goal of the controller is to process the robots in a way that maximizes the profit gained from all robots, throughout the operation of the system. This yields an instance of SWV, where job j has a processing time p_j, and the window $a_j = p_j + \ell_j$.

Commercial Broadcast: In a broadcasting system which transmits data, e.g., commercials on a running banner, there is a profit associated with the transmission of each commercial. This profit is gained only if some predefined period of time elapsed since the previous transmission. The goal is to broadcast the commercials in a way that maximizes the overall profit of the system. Thus, we get an instance of SWV, where the window of each job corresponds to the time interval between consecutive transmissions of each commercial.

1.1 Our Results and Related Work

The problem of scheduling with vacations belongs to the class of *periodic scheduling* problems where each job may be scheduled infinite number of times. Periodic scheduling is a well-studied problem. Problems of this type arise in many areas, including production planning, machine maintenance, telecommunication systems, media-on-demand, image and speech processing, robot control/navigation systems, and databases query processing. In the paper [12], which introduced the periodic scheduling problem, the goal is to schedule each of the jobs such that the average gap between two executions of job j is a_j using the minimum number of machines. Another early example is the *chairman assignment* problem [13], in which the number of executions of job j in any prefix of the schedule of length t must be at least $\lfloor t/a_j \rfloor$ and at most $\lceil t/a_j \rceil$. For both problems the *earliest deadline first* algorithm was shown to be optimal.

Our problem is different from previously studied variants of periodic scheduling in two major aspects. First, our goal is to maximize the total profit of the server. This may result in lack of fairness. Indeed, in a schedule which maximizes the profit, some of jobs may be waiting forever, while other (more profitable) jobs are repeatedly processed as soon as they return from vacation. In other works (e.g., [12, 7, 8]), the performance of a periodic schedule is measured by some fairness criterion, or (e.g. [2, 9]) the two objectives are combined; that is, each job comprises a mandatory and an optional part, with which a non-decreasing reward

2

function is associated. In other variants of the periodic scheduling problem (see, e.g., in [1]), the processing of each job incurs a service cost, depending on the time that elapsed since the last service of this job, and the goal is to minimize the cost of the schedule, however, the cost increases with job delays − implying that an optimal solution requires some fairness.

The other major difference from well-studied variants of periodic scheduling is that in our problem, due to the vacation requirement, jobs cannot be processed too often. In the *windows scheduling* problem [4, 5], the parameter associated with each job gives the *maximal* gap between any two executions of job j, however, in any feasible schedule a job may be processed more often. In the *periodic machine scheduling problem (PMSP)* [3], this parameter specifies the *exact* gap between any two executions of j. To the best of our knowledge, there is no earlier work in which only the *minimal* gap between any two executions of a job is given.

We first show that SWV is NP-hard already for jobs with unit length and unit profits. Since the total number of different configurations of the system is finite, an optimal solution can be found (inefficiently) by applying the known *buffer scheme* designed for the windows scheduling problem [6]. We present several approximation algorithms for various subclasses of instances. A simple greedy algorithm yields a 2-approximation for preemptive scheduling of jobs with arbitrary processing times and arbitrary profits. For the special case of unit processing times, we present a 1.66-approximation algorithm, for instances with arbitrary profits, and a 1.39-approximation algorithm for instances with the same (unit) profits. For the latter case, we also show that when the load generated by an instance is sufficiently large (in terms of n and m), any algorithm that uses no intended idle times yields an optimal schedule.

Our approximation algorithms for unit length jobs apply a transformation of the instance to an *aligned instance*, where all window sizes satisfy certain properties. Our approximation technique, based on finding an optimal schedule for the resulting aligned instance, builds on a technique of [5] for the windows scheduling problem. To obtain better approximation ratio, we extend this technique: Our algorithm for unit processing times and unit profits applies the *best align* technique, which finds an aligned instance that yields the best approximation ratio for the original instance.

2 Preliminaries

2.1 Definitions and Notation

Denote by a_j the *scheduling window* (for short, *window*) of job j. The *load* incurred by job j is $\ell(j) = p_j/a_j$. The best service job j can receive is to be processed whenever it returns from vacation (i.e., no waiting time). Thus, the load of job j represents the average processing requirement of j per time unit. The total load of the input is $L(J) = \sum_{j \in J} \ell(j)$. In the special case of unit processing times (i.e., for all j, $p_j = 1$), we get that $\ell(j) = 1/a_j$, and $L(J) = \sum_{j \in J} 1/a_j$. Each job is associated with an additional parameter, b_j, the profit gained from each execution of job j.

We say that a job j is *heavier* than j' if $b_j/p_j > b_{j'}/p_{j'}$, or $b_j/p_j = b_{j'}/p_{j'}$ and $j > j'$. A heavier job achieves larger profit per unit time of its execution; the second part of the definition

only breaks the ties consistently for the whole instance. Let $J_{\geq j}$ denote the set of jobs at least as heavy as j.

To define the profit of a schedule formally, let $compl_T^S(j)$ denote the number of times job j is completed in the first T time units in schedule S. The profit of schedule S is defined as

$$profit(S) = \liminf_{T \to \infty} \frac{1}{T} \sum_{j=1}^{n} compl_T^S(j) b_j.$$

A schedule is periodic with period p if for any time T, the jobs scheduled at time T are the same as the jobs scheduled at time $T + p$. It is easy to see that there always exists a periodic optimal schedule: Define a configuration at time T to consist of the information of how much of each job is currently executed, or how much time has elapsed since its last completion. An optimal schedule can always be chosen so that for two times T and T' with the same configurations, the remainder of the schedule is the same. Since the number of configurations is finite, such a schedule is periodic.

In the remainder of paper we consider only periodic schedules. For a schedule S with period p, in the definition of profit, $\liminf$ can be replaced by $\lim$ which always exists and it is actually equal to the value of the remaining expression for $T = p$. Similarly, it is meaningful to speak about the load (relative frequency) of a job j in S; we denote it by $\ell^S(j)$. Finally, $L^S(J') = \sum_{j \in J'} \ell^S(j)$ denotes the load of a set of jobs $J' \subseteq J$ in a schedule S.

The next lemma will be useful in comparing our schedules to an optimal schedule S^*.

Lemma 2.1 *Let S^* be an arbitrary schedule for an instance J and let $R \geq 1$.*

(i) *Suppose that for a schedule S for J, and for any job j, $L^S(J_{\geq j}) \geq L^{S^*}(J_{\geq j})/R$. Then $profit(S) \geq profit(S^*)/R$.*

(ii) *Suppose that for a distribution of schedules $\mathcal{S}$ for J, and for any job j, $\mathbf{Exp}_{S \in \mathcal{S}}[L^S(J_{\geq j})] \geq L^{S^*}(J_{\geq j})/R$. Then one of the schedules $S \in \mathcal{S}$ satisfies $profit(S) \geq profit(S^*)/R$.*

Proof: The lemma follows from the fact that the profit of any schedule S can be computed as $\int_{\beta=0}^{\infty} L^S(J_\beta) \, d\beta$, where $J_\beta = \{j \in J \mid b_j \geq \beta\}$. ∎

2.2 Hardness Results

The hardness of SWV with a single machine, unit profits, and unit processing times is shown by a reduction from the *periodic machine scheduling problem (PMSP)*, which is known to be NP-hard (see in [3]). In PMSP, there is a set of n machines; the j-th machine requires maintenance every a_j time units, where $\sum_{j=1}^{n} 1/a_j \leq 1$. The maintenance of any machine takes one time unit, and at any time only one machine can be maintained. It is assumed that the first maintenance of machine j can be done in any of the first a_j time slots. The goal is to find a maintenance schedule that satisfies *exactly* all the requirements, that is, for all j, machine j is maintained exactly once in any window of a_j time slots.

4

Theorem 2.2 *SWV is NP-hard, even with a single machine, unit profits and unit processing times.*

Proof: Given an instance of PMSP with n machines in which the j-th machine requires maintenance every a_j time units, construct an instance of SWV with unit profits, unit processing times, and scheduling windows a_j for jobs $j = 1, \ldots, n$. Since $\sum_{j=1}^{n} 1/a_j \leq 1$, there is a feasible schedule for PMSP if and only if the average profit per time-unit from job j is exactly $1/a_j$, that is, there is a schedule of the SWV instance with average profit $\sum_{j=1}^{n} 1/a_j$. ∎

Remark 2.1 When $L(J) \geq 1$, the problem is still NP-hard for unit profits and unit processing times. The proof is similar to the hardness proof for the windows scheduling problem given in [5]. Given an instance with $L(J) < 1$ it can be extended into a one with $L(J') = 1$ in a way that does not affect the schedule of the original jobs. Let $A = LCM(a_1, a_2, ..., a_n)$, that is, each of the original windows divides A. Add to J dummy jobs each having window A, such that the total load is 1. Then a schedule of the resulting instance with average profit 1 induces a schedule of the original instance with average profit $\sum_{j=1}^{n} 1/a_j$, and vice versa. For arbitrary profits a single dummy job with $b_j = \varepsilon$ and $a_j = 1$ is sufficient.

2.3 An Optimal Super-Polynomial Algorithm

The general buffer scheme [6] is a tool designed for solving optimally the windows scheduling problem. We now explain how this tool can be adjusted to solve optimally the problem of scheduling with vacations. We explain this adjustment for unit-length jobs and arbitrary profit. It can be extended to handle jobs with arbitrary lengths similar to the extension for windows scheduling shown in [6]. The buffer scheme is based on representing the system as a nondeterministic finite state machine in which any directed cycle corresponds to a valid periodic schedule and vice-versa. Let $a^* = \max_j\{a_j\}$. The state of a job is represented using a set of buffers, $B_0, B_2, \ldots, B_{a^*-1}$. Each job is stored in some buffer. A job is stored in B_i when i time slots remain to finish its vacation. Initially, all jobs are ready to be processed so they are all stored in B_0. In each iteration, the scheme selects for scheduling at most m jobs from B_0 and updates the content of the buffers:

- For all $i > 0$, all jobs located in B_i are moved to B_{i-1} (note that none of them is scheduled).

- Each scheduled job, j, is moved from B_0 to B_{a_j-1} (this ensures that at least $a_j - 1$ time slots will pass before j is selected again).

Note that the total number of buffer configurations is at most $\Pi_j a_j$. In the state machine, each configuration is represented by a vertex, and there is an edge connecting configuration a to configuration b if it is possible to move from the state corresponding to a to the state corresponding to b in a single time slot. Each edge $e = (a, b)$ has a weight b_e whose value is the profit gained by moving from a to b, that is, the total profit of the jobs selected for execution.

A periodic schedule with cycle t corresponds to a cycle of length t in the state machine. The profit of this schedule is given by the average edge weight along this cycle. Thus, an optimal schedule corresponds to a cycle with maximal mean weight. Such a cycle can be found by using the classic *Max Cycle Mean* algorithm of Karp [10]. The running time is polynomial in the graph size, which is polynomial in $\Pi_j a_j$. Clearly, this algorithm is applicable only for small instances, however, since in general the problem is NP-hard, we cannot expect efficient optimal solutions.

3   A 2-approximation Algorithm

We describe below an algorithm for preemptive scheduling of arbitrary-length jobs. Note that in the preemptive scheduling model, the execution of a preempted job can be resumed any time. The vacation-constraint refers only to the gaps between *different* executions of a job. Recall that m is the number of available machines.

Algorithm GREEDY

At any time t, schedule for one time unit the m heaviest available jobs.

It is easy to see that GREEDY always generates a periodic schedule, as the number of different states is finite.

Theorem 3.1 GREEDY *is a 2-approximation algorithm.*

Proof: Let S be the schedule produced by greedy and $R = 2$. Order the jobs from the heaviest one, and denote them $j_1, j_2, \ldots, j_n$ in this order.

We prove that for every i,

$$L^S(J_{\geq j_i}) \geq \min\{m/2, L(J_{\geq j_i})/2\}.$$

Since the optimal schedule S^* satisfies $L^{S^*}(J') \leq \min\{m, L(J')\}$ for any $J' \subseteq J$, the theorem then follows from Lemma 2.1.

To prove the claim, we distinguish between two cases.

First, assume that for some i, $\ell^S(j_i) \leq \ell(j_i)/2$. If two consecutive completion times of the job are $a_{j_i} + T$ time units apart then between the two completion times there are at least T time slots in which j_i was available but not scheduled; by the definition of the algorithm, heavier jobs were scheduled in these time slots on all machines. Consequently, the case assumption implies that in half of the slots of the schedule jobs heavier than i were executed. Thus, $L^S(J_{\geq j_i}) \geq m/2$, and the same follows also for any $i' \geq i$.

Second, by induction on i we prove that until the first case occurs,

$$L^S(J_{\geq j_i}) \geq L(J_{\geq j_i})/2.$$

Since the first case is excluded, we have $\ell^S(j_i) > \ell(j_i)/2$ and the inductive step follows by summing this inequality and the induction assumption for $i - 1$. (If $i = 0$ then the claim follows trivially.) ∎

We conjecture that our analysis of GREEDY is not tight and in fact its approximation ratio is $e/(e-1) \approx 1.58$. The following is an example of an $e/(e-1)$ ratio for unit-length jobs: Let a be an integer and let $A = \{a + 1, a + 2, \ldots, ea\}$. For sufficiently large a it holds that $\sum_{i \in A} 1/i = \ln ea - \ln a = \ln e = 1$. Let $m = lcm\{a + 1, a + 2, \ldots, ea\}$. That is, each of the numbers $a+1, a+2, \ldots, ea$ divides m. The instance consists of m machines, and m jobs having vacation i for each $i = a, a + 1, \ldots, ea - 1$, a total of $(e-1)am$ jobs. The profits are almost identical with a slight preference to jobs having long vacation.

The GREEDY schedule: In the first slot the most profitable jobs are those with vacation $ea - 1$, so these m jobs are scheduled first on each of the m machines. Next (on all m machines) are the jobs with vacation $ea - 2$ and so on, on each of the m machines the resulting schedule is $[ea - 1, ea - 2, ..., a, *, \ldots, *]$ repeated infinitely, where $*$ denotes idle. Since all jobs are on vacation there are a idle slots after each sequence of $ea - a$ busy slots. The average profit per slot on each of the machines is therefore $((ea - a)/ea) - \varepsilon = ((e-1)/e) - \varepsilon$ where ε is a constant depending on the (small) jobs' profit difference.

An optimal schedule: m/ea machines schedule with no idle times all jobs having vacation $em - 1$. In general, m/i machines schedule with no idle the jobs having vacation $i - 1$. Since $\sum_{i \in A} 1/i = 1$, we have enough jobs to 'saturate' this way all the machines with no idle times, and therefore the average profit per slot is $1 - \varepsilon'$. The ratio between the optimal profit and GREEDY's profit tends to $e/e - 1$.

4   Unit Processing Times

Call an instance *aligned* if there exists an integer q such that for each j, $a_j = q2^{\alpha_j}$ for some integer $\alpha_j \geq 0$, or $a_j = 1$. (Note that this generalizes the case when all the windows are powers of 2.) The following result is due to [5].

Theorem 4.1 *An optimal solution for an aligned instance can be computed in polynomial time.*

For the sake of completeness, we describe such an optimal solution. Schedule the jobs in order of increasing windows, always keeping the period of the schedule to be equal to the maximal window processed so far. The machines are utilized in an arbitrary order. When a new job is processed, if needed, increase the period by repeated doubling and repeating the current schedule, until the period is equal to the currently processed window. Then if possible, schedule the current job in an empty slot on some machine. Otherwise, if the current job is heavier than some of the scheduled jobs, put this job in a slot of the lightest scheduled job. (This light job can still be scheduled is some of the slots, due to a possible previous doubling of the period.) Note also that this algorithm works also if we start not with an empty schedule but with some schedule with period q where some slots are already used to run some other jobs.

An alternative view of this schedule S will be useful. Let j be the heaviest job such that $L(J_{\geq j}) \geq m$. Then the schedule above has the property that all jobs j' heavier than j are scheduled at their maximal rate, i.e., $\ell^S(j') = \ell(j')$, and no jobs lighter than j occur in the schedule. The borderline job j has such a rate that all the rates sum to 1.

It follows that a simple 2-approximation algorithm can be achieved by rounding the window of each job to the next higher power of 2 and using the optimal algorithm for aligned instances. Below we give two algorithms that refine this idea and achieve a better approximation ratio. These algorithms are deterministic: randomization is used only for their analysis.

4.1 A 1.66-approximation Algorithm for Arbitrary Profits

Algorithm SIMPLEALIGN

Produce 2 instances J' and J''. The instance J' has all the windows rounded up to the next power of 2, the instance J'' has all the windows rounded up to the next number of the form $3 \cdot 2^\alpha$ for some integer $\alpha \geq 0$, with the exception of jobs with window 1 which remain unchanged. Produce optimal schedules S' and S'' for J' and J'' and choose the better one.

Theorem 4.2 SIMPLEALIGN *is a 5/3-approximation algorithm for unit processing times and arbitrary profits.*

Proof: Consider a distribution which chooses S' with probability 2/5 and S'' with probability 3/5. We prove that the condition of Lemma 2.1 is satisfied with $R = 5/3$. We use the following claim, where by $\ell^{J'}(j)$ we denote the load of a job j rounded as in the instance J', and similarly for $\ell^{J''}$ and J''.

Claim 4.3 *For any job j, $\frac{2}{5}\ell^{J'}(j) + \frac{3}{5}\ell^{J''}(j) \geq \ell(j)/R$.*

Proof: We consider two cases:

(i) If for some integer α, $2^{\alpha+1} \leq a_j \leq 3 \cdot 2^\alpha$, then

$$\frac{2}{5}\ell^{J'}(j) + \frac{3}{5}\ell^{J''}(j) \geq \frac{2}{5} \cdot \frac{1}{2^{\alpha+2}} + \frac{3}{5} \cdot \frac{1}{3 \cdot 2^\alpha} = \frac{3}{5 \cdot 2^{\alpha+1}} \geq \frac{\ell(j)}{R}.$$

(ii) Otherwise, for some integer α, $3 \cdot 2^\alpha < a_j < 2^{\alpha+2}$, in which case

$$\frac{2}{5}\ell^{J'}(j) + \frac{3}{5}\ell^{J''}(j) \geq \frac{2}{5} \cdot \frac{1}{2^{\alpha+2}} + \frac{3}{5} \cdot \frac{1}{3 \cdot 2^{\alpha+1}} = \frac{1}{5 \cdot 2^\alpha} \geq \frac{\ell(j)}{R}.$$

∎

Now, given an optimal schedule S^* and a job j, we prove the assumption of Lemma 2.1. We distinguish between two cases.

(a) First, assume that both $L^{S'}(J_{\geq j}) < m$ and $L^{S''}(J_{\geq j}) < m$. Then $L^{S'}(J_{\geq j}) = \sum_{j \in J_{\geq j}} \ell^{J'}(j)$, by the construction of an optimal schedule for aligned instances, and similarly for S''. From Claim 4.3, we obtain $\mathbf{Exp}[L^S(J_{\geq j})] \geq L(J_{\geq j})/R \geq L^{S^*}(J_{\geq j})/R$.

8

(b) Otherwise it must be the case that $L^{S'}(J_{\geq j}) = m$ or $L^{S''}(J_{\geq j}) = m$ (note that the load in a schedule cannot be larger than m). Either way, it implies that $L(J_{\geq j}) \geq m$ in the original instance. In both J' and J'', the size of each window is at most doubled, therefore $L^{S'}(J_{\geq j}) \geq m/2$ and $L^{S''}(J_{\geq j}) \geq m/2$. Then the average load is bounded by $\mathbf{Exp}[L^S(J_{\geq j})] \geq 2/5 + 3/5 \cdot m/2 = 7m/10 \geq m/R \geq L^{S^*}(J_{\geq j})/R$.

The proof is completed by an application of Lemma 2.1. ∎

4.2 A 1.39-approximation Algorithm for Unit Profits

Now we focus on the special case where for each job $b_j = 1$. Thus, the average profit is equal to the load of the schedule. In particular, if the schedule has no idle time, it is optimal.

Instead of using two schedules with periods 2 and 3 times a power of 2, the next algorithm uses one of k schedules with periods $q = k + 1, k + 2, \ldots, 2k$ times a power of 2. For a large but constant k, the approximation ratio approaches $2 \ln 2 \approx 1.39$.

The jobs with large windows are rounded similarly as in the previous proof. The jobs with small windows are handled separately. Given a constant q and a set of jobs with windows at most q, we find an optimal schedule with period q. If a number of machines is constant, then the number of such schedules is constant and we can find an optimal one simply by exhaustive search. For m arbitrary, we use Lenstra's polynomial algorithm for integer programming in fixed dimension, see [11], similarly as in its other applications in scheduling. Since q is a constant, we have a constant number of job types specified by their windows, and for each type of jobs we have a constant number of patterns specifying in which time slots it is run. The variables of the integer program correspond to the number of jobs of each type following each pattern; the number of these variables is a constant exponential in q. The constraints specify that (i) the number of scheduled jobs of each type equals the number of such jobs in the instance and that (ii) the number of jobs scheduled in a given time unit is at most m. Feasible integral solutions then correspond to schedules in a straightforward way. Given an instance, the integer program can be produced in linear time, and solved in time polylogarithmic in n and m (with a multiplicative constant doubly exponential in q).

> **Algorithm** BESTALIGN
>
> Let $K = \lceil 1/\varepsilon \rceil$. For all values $k \in \{K, 4K, \ldots, 4^x K, \ldots, 4^K K\}$ and for all values $q = \{k + 1, k + 2, \ldots, 2k\}$ generate a schedule as follows.
>
> Let J' be the set of jobs with windows at most $k/4$. Let J'' be an instance of jobs with windows larger than k. (Note that J' and J'' depend only on k.) Let J''' be an aligned instance obtained from J'' so that the window of each job is rounded up to the next higher number of the form $q2^\alpha$ for integer $\alpha \geq 0$.
>
> Find the best schedule with period q for J'; since q is a constant, this can be done as described in the previous paragraph. Then schedule J''' in the remaining slots using the optimal algorithm for aligned instances.
>
> Output the best schedule over all values of k and q.

Theorem 4.4 BESTALIGN *is* $(2\ln 2 + O(\varepsilon))$-*approximation algorithm for the jobs with unit processing times and unit profits. Thus there exist an R-approximation algorithm for any* $R > 2\ln 2 \approx 1.386$.

Proof: Let S^* be an optimal schedule for a given instance.

For the proof, fix a value of k among those used in the algorithm so that the contribution of the jobs with windows between $k/4$ and k to $L(S^*)$ is at most $\varepsilon \cdot L(S^*)$. Since there are $\lceil 1/\varepsilon \rceil$ choices of k, one of them has a sufficiently small contribution. This defines the sets of jobs J' and J''. Let S' and S'' denote the schedule S^* restricted to J' and J'', respectively, i.e., all the other jobs are simply removed from the schedule.

For any q, let S_q be the schedule produced by the algorithm for this choice of q and for k fixed as above. If any of S_q has load m, it is optimal for J and the theorem follows. Thus, for the remaining proof we can assume that S_q always schedules all the jobs in J''' with their maximal load, i.e., $L^{S_q}(J''') = L(J''')$.

We prove that one of the schedules S_q has load at least $3/4 \cdot L^{S'}(J') + L(J'')/(2\ln 2 + O(\varepsilon))$. The theorem then follows since by the choice of k, $L^{S^*}(J) \leq (1 + O(\varepsilon))(L^{S'}(J') + L^{S''}(J'')$, furthermore $L^{S''}(J'') \leq L(J'')$, and $2\ln 2 > 4/3$.

The schedule produced by BESTALIGN for J' has load at least $3/4$ of $L(S')$: If in S' we take a segment of length $q - k/4$ with the maximal load and append to it $k/4$ empty slots, we get a schedule with period q and load at least $3/4$ of $L(S')$. The optimal schedule for J' with period q chosen by BESTALIGN has at least the same load.

Now we complete the proof of theorem by showing that for some distribution over schedules S_q, $\mathbf{Exp}[L^{S_q}(J''')] \geq L(J'')/(2\ln 2 + O(\varepsilon))$. Define the probability distribution so that the probability of choosing S_q is

$$\pi_q = \frac{X}{q-1}. \tag{1}$$

where X is chosen so that the sum of the probabilities is 1, i.e.,

$$\left(\frac{1}{k} + \ldots + \frac{1}{2k-1} \right) X = 1. \tag{2}$$

We proceed with the proof job by job, similarly as in Theorem 4.2. Since S_q schedules all the jobs with maximal rate, it is sufficient to prove for each job j that its expected load in J''' (i.e., after rounding) is at least $\ell(j)/(2\ln 2 + O(\varepsilon))$. Assume that integers $\alpha \geq 0$ and $t \in \{k, \ldots, 2k-1\}$ satisfy $t2^\alpha < a_j \leq (t+1)2^\alpha$ (for each j, there exist unique values of t and α). Then the average load of j in J''' is

$$
\begin{aligned}
\sum_{q=k+1}^{t} \pi_q \cdot \frac{1}{q2^{\alpha+1}} + \sum_{q=t+1}^{2k} \pi_q \cdot \frac{1}{q2^\alpha} &= t \left(\sum_{q=k+1}^{t} \pi_q \cdot \frac{1}{2q} + \sum_{q=t+1}^{2k} \pi_q \cdot \frac{1}{q} \right) \frac{1}{t2^\alpha} \\
&\geq t \left(\sum_{q=k+1}^{t} \pi_q \cdot \frac{1}{2q} + \sum_{q=t+1}^{2k} \pi_q \cdot \frac{1}{q} \right) \ell(j).
\end{aligned}
$$

It remains to bound the coefficient on the right-hand side. From (1), we have

$$\frac{\pi_q}{q} = \frac{X}{(q-1)q} = \left(\frac{1}{q-1} - \frac{1}{q}\right) X,$$

and substituting this and telescoping the sums we have

$$t\left(\sum_{q=k+1}^{t} \pi_q \cdot \frac{1}{2q} + \sum_{q=t+1}^{2k} \pi_q \cdot \frac{1}{q}\right) = tX\left(\frac{1}{2k} - \frac{1}{2t} + \frac{1}{t} - \frac{1}{2k}\right) = \frac{X}{2}.$$

Using (2), we have

$$\frac{2}{X} = 2\left(\frac{1}{k} + \ldots + \frac{1}{2k-1}\right).$$

$$\leq 2\int_k^{2k} \frac{dx}{x} + \frac{2}{k} = 2(\ln(2k) - \ln k) + \frac{2}{k} \leq 2\ln 2 + O(\varepsilon).$$

This completes the proof of the theorem. ∎

4.3 Instances with a Large Load

In the case where $p_j = b_j = 1$, the goal is to maximize the utilization of the machine. A simple observation is that the best schedule one can expect is one in which the machines are busy all the time. In this subsection we focus on two cases where such a schedule can be generated.

Our analysis of SIMPLEALIGN shows that if $L(J) \geq 5m/3$, then SIMPLEALIGN produces a schedule with load m: One of the schedules S' and S'' is guaranteed to have at least this load, and no larger load is possible.

The next theorem gives a guarantee when even any greedy algorithm produces a schedule with full utilization of all machines. In other words, no machine is never idle, regardless of the jobs selected to be scheduled at any time slots.

Theorem 4.5 *Suppose that $n = km+r$ for some $r < m$. (I.e., $k = \lfloor n/m \rfloor$ and $r = n \mod m$.)*

If $L(J) > mH_k - 1 + \frac{r+1}{k+1}$ then any algorithm with no intended idle times achieves the optimal profit. In particular, for $m = 1$, the condition is $L(J) > H_{n+1} - 1$.

Proof: We show that if a machine is idle, $L(J)$ must be small. Assume that some machine is idle for the first time at time t. Thus, none of the n jobs is available, meaning that all jobs are processed on the other machines or on vacation. Order the jobs in the order of the last execution up to time t, including possible executions on the other $m-1$ machines at time t. The last $m-1$ jobs in this order are possibly executed at time t, we have no bound on their vacation, so each can contribute as much as 1 to $L(J)$. The previous m jobs in this order are scheduled no later than at time $t-1$, they are on vacation at time t, thus they have $a_j \geq 2$ and contribute at most $1/2$ each to $L(J)$. Similarly, the previous m jobs are scheduled no later

than at time $t - 2$ and contribute at most $1/3$ each to $L(J)$, and so on. The last $r + 1$ jobs contribute at most $1/(k + 1)$ each. The total load of the input is at most

$$L(J) \leq m - 1 + m \left(\frac{1}{2} + \frac{1}{3} + \ldots + \frac{1}{k} \right) + \frac{r + 1}{k + 1} = mH_k - 1 + \frac{r + 1}{k + 1}.$$

Thus $L(J) > mH_k - 1 + \frac{r+1}{k+1}$ implies that there is no idle time. ∎

Acknowledgment We thank Gerhard Woeginger for stimulating discussions on this paper.

References

[1] S. Anily, J. Bramel, Periodic scheduling with service constraints. *Operations Research*, Vol. 48, pp. 635-645, 2000.

[2] H. Aydin, R. Melhem, D. Mosse, P. Mejia-Alvarez, Optimal Reward-Based Scheduling of Periodic Real-Time Tasks. In 20th IEEE Real-Time Systems Symposium, 1999.

[3] A. Bar-Noy, R. Bhatia, J. Naor and B. Schieber, Minimizing Service and Operation Costs of Periodic Scheduling. In *Proc. of SODA*, 1998.

[4] A. Bar-Noy and R. E. Ladner. Windows Scheduling Problems for Broadcast Systems. In *Proc. of SODA*, 2002.

[5] A. Bar-Noy, R. E. Ladner and T. Tamir. Windows Scheduling as a Restricted Version of Bin Packing. In *Proc. of SODA*, 2004.

[6] A. Bar-Noy, R. E. Ladner and T. Tamir. A General Buffer Scheme for the Windows Scheduling Problem. In *Proc. of WEA*, 2005.

[7] S.K. Baruah, N.K. Cohen, C.G. Plaxton, D.A. Varvel, Proportionate Progress: A Notion of Fairness in Resource Allocation. *Algorithmica*, vol.15 (6), pages 600 –625, 1996.

[8] S. K. Baruah, S-S. Lin. Pfair Scheduling of Generalized Pinwheel Task Systems. *IEEE Trans. on Comp.*, Vol. 47, 812–816, 1998.

[9] J. Chung, J.W.S. Liu, K. Lin. Scheduling Periodic Jobs that Allow Imprecise Results. *IEEE Trans. on Comp.*, Vol. 39 (9),pp. 1156–1174, 1990.

[10] R. M. Karp. A Characterization of the Minimum Cycle Mean in a Digraph. *Discrete Math.*, 23:309-311, 1978.

[11] H.W. Lenstra. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, Vol. 8, 538–548, 1983.

[12] C. L. Liu and J. W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *Journal of the ACM*, Vol. 20, No. 1, 46-61, 1973.

[13] R. Tijdeman. The Chairman assignment Problem. *Discrete Mathematics*, Vol. 32, 323-330, 1980.