Temporary Tasks Assignment Resolved

Amitai Armon * Yossi Azar † Leah Epstein ‡ Oded Regev §

Abstract

Among all basic on-line load balancing problems, the only unresolved problem was load balancing of temporary tasks on unrelated machines. This open problem exists for almost a decade, see [Borodin El-Yaniv]. We resolve this problem by providing an unapproximability result. In addition, a newer open question is to identify the dependency of the competitive ratio on the durations of jobs in the case where durations are known. We resolve this problem by characterizing this dependency. Finally, we provide a $PTAS$ for the off-line problem with a fixed number of machines and show a 2 unapproximability for the general case.

1 Introduction

On-line load balancing was extensively studied in the last decade (e.g., [1, 2, 3, 5, 10, 11, 13, 16, 18, 20, 23]). The basic problem contains the identical, related, restricted and unrelated models for permanent and temporary tasks. Tight bounds were given to all these problems except one: the assignment of temporary tasks to unrelated machines remained open. We present an unapproximability result by employing a cyclic load transfer method. Another more recent open question posed in [12] is whether the case where job durations are known at their arrival is provably harder than that of permanent jobs. We answer this question in the affirmative. We first summarize the results presented in this paper. Let m denote the number of machines and T the duration of the longest job.

- For on-line unrelated assignment of temporary tasks with unknown durations, we show an $\Omega(m/\log m)$ lower bound which almost matches a trivial $O(m)$ upper bound. For randomized algorithms we show an $\Omega(m/(\log m)^2)$ lower bound.
- For on-line restricted assignment of temporary tasks with known durations, we present a lower bound of $\Omega(\sqrt{m})$ and of $\Omega(\sqrt{\frac{\log T}{\log\log T}})$ on the competitive ratio of any on-line algorithm (deterministic or randomized). These lower bounds also hold for assignment on unrelated machines.
- For offline assignment of temporary tasks on unrelated machines, we present a PTAS for the case where the number of machines is fixed. For the case where the number of machines is not fixed, we present a lower bound of 2 (provided that $P \neq NP$).

We also provide two additional results for interesting special cases of temporary tasks assignment in the unrelated machines model with unknown durations. Specifically, we show tight results for certain cases of the related-restricted model and of the restricted machines model.

Definitions and previous results: We consider the problem of non-preemptive load balancing of temporary tasks on m unrelated machines. Each job (task) has an arrival time and a departure time, and should be assigned to one machine immediately upon its arrival. Each job j is associated with a loads vector $(w_j(1), ..., w_j(m))$. If job j is assigned to machine i, it increases the load of machine i by its *weight* on that machine, $w_j(i)$, for the duration of the job. The load on a machine at a certain time is the sum of the loads caused by the jobs assigned to it at that time. The goal is to minimize the maximum load over machines and time. Note that the load and the time are two separate axes of the problem. In the *known durations* setting we assume that when a job arrives its duration is given to the on-line algorithm and in the *unknown durations* setting we assume that the departure time is known only when the job actually departs.

An important special case of the unrelated machines model is the special case called the *restricted assignment model.* In this case, a job j can only be assigned to a subset of the machines that depends on the job. On each of these machines it causes the same load w_j. This is

*Department of Computer Science, Tel Aviv University, Tel-Aviv, 69978, Israel. E-Mail: armon@tau.ac.il.

†Department of Computer Science, Tel Aviv University, Tel-Aviv, 69978, Israel. E-Mail: azar@tau.ac.il. Research supported in part by the Israel Science Foundation and by a grant of the European Commission.

‡School of Computer and Media Sciences, The Interdisciplinary Center, Herzliya, Israel. E-Mail: epstein.leah@idc.ac.il.

§Institute for Advanced Study, Princeton, NJ. E-Mail: odedr@ias.edu. Research supported in part by NSF grant CCR-9987845

equivalent to a loads vector that contains only the values w_j or ∞. The results in [4] (and later in [21]) show a lower bound of $\Omega(\sqrt{m})$ on the competitive ratio that any on-line algorithm (deterministic or randomized) may have for restricted assignment of temporary tasks with unknown durations. An on-line algorithm for this problem with a competitive ratio of $O(\sqrt{m})$, was later presented in [6] (the "Robin-Hood" algorithm) thereby proving that the lower bound of $\Omega(\sqrt{m})$ is tight.

Unknown durations: Apart from a trivial upper bound, no on-line algorithm for the unrelated assignment model exists. The $\Omega(\sqrt{m})$ lower bound for restricted assignment mentioned above ([4]) dates back to 1992 but was still the best lower bound for unrelated assignment. Our results show that a trivial $O(m)$ competitive algorithm is almost optimal hence proving the unapproximability of this model. Specifically, by using a cyclic load transfer method we achieve an $\Omega(m/\log m)$ lower bound. The open problem mentioned in [12] regarding the existence of a better approximation algorithm is thus answered negatively. We extend the unapproximability result to randomized algorithms as well.

Known durations: The competitive ratio of $\Theta(\sqrt{m})$ for the restricted assignment model is usually regarded as a high competitive ratio, so in [6] it was suggested to consider this problem in the known-durations case, hoping to beat the $\Omega(\sqrt{m})$ bound, and get a polylogarithmic competitive ratio. This seemed plausible since the best known lower bound for the known durations case was only $\Omega(\log m)$ [7] (proven for the special case of permanent tasks). Indeed, the work of [6] made a step in this direction by showing an $O(\log mT)$ - competitive algorithm, where T is the duration of the longest job (in discrete time units). The competitive ratio of this algorithm is lower than $\Theta(\sqrt{m})$ for a large range of T. The obvious intriguing question, also presented in the survey of [3], was whether we can indeed improve on the $\Omega(\sqrt{m})$ lower bound and achieve a $O(\mathit{polylog}(m))$-competitive algorithm. Surprisingly, we answer this question negatively, by extending the lower bound of [21] to a lower bound of $\Omega(\sqrt{m})$ that holds for the known durations case. This bound is tight by the upper bound for unknown durations. Our $\Omega(\sqrt{m})$ lower bound holds for unrelated machines as well (since restricted assignment is a special case). This result also answers an open question presented by Borodin and El-Yaniv in [12]. They asked whether there is any machine model in which one can prove a lower bound for temporary tasks assignment with known durations which is higher than the competitive ratio of permanent tasks assignment in that model. Our $\Omega(\sqrt{m})$ lower bound for restricted assignment of temporary tasks with known durations is strictly higher than the $O(\log m)$ competitive ratio for the case of permanent tasks. We note that as a function of T our lower bound is $\Omega(\sqrt{\frac{\log T}{\log\log T}})$ which can be compared to the $O(\log T)$ upper bound of [6] (assuming $T \geq m^{\epsilon}$).

Offline results: Next we consider the assignment of temporary tasks on unrelated machines in the offline setting. In the special case of *permanent* tasks, the tasks do not depart (all the departure times equal ∞). Horowitz and Sahni presented an FPTAS for permanent tasks assignment on a fixed number of unrelated machines (i.e., the number of machines is not a part of the input) [17]. A PTAS for this problem was also presented by Lenstra et al. [19]. For permanent tasks assignment on an arbitrary number of unrelated machines, Lenstra et al. [19] and Shmoys and Tardos [22] presented algorithms with an approximation ratio of 2. In addition, Lenstra et al. proved that no algorithm can reach an approximation-ratio better than $\frac{3}{2}$ for the arbitrary number of machines case, unless $P = NP$ [19].

Unlike the permanent case, solving the problem of *temporary* tasks assignment on unrelated machines cannot be done by standard rounding techniques. The problem arises from the two separate axes of the problem and this extra dimension is known to turn problems into intractable ones or very hard to approximate [15, 14, 9]. In order to obtain a PTAS after all, we had to use a two-dimensional rounding technique. While the above applies to the case of fixed number of machines, we also prove that when the number of machines is not fixed (i.e. is part of the input), no algorithm can achieve an approximation ratio lower than 2 unless $P = NP$. This lower bound is higher than the $\frac{3}{2}$ lower bound for permanent tasks and almost separates the two settings.

The problem of temporary tasks assignment on identical machines is another special case of our problem which was considered in [8] and provides the foundation for our algorithm. In this special case, the load that a job causes depends only on the job and it is identical for all the machines (i.e. $w_j(i) = w_j$ for $1 \leq i \leq m$). A PTAS for this problem in the case where the number of machines is fixed was presented in [8]. They also proved a lower bound of $\frac{3}{2}$ for the case of an arbitrary number of machines, provided that $P \neq NP$. Again, the lower bound that we present for our problem is higher than the lower bound for this special case.

2 Tasks of Unknown Duration

2.1 Unapproximability of the Unrelated Model

In this section we present the unapproximability result for online load balancing of temporary tasks on unrelated machines. Namely, we show a lower bound of

$\Omega(m/\log m)$ almost matching an upper bound of $O(m)$. It can be seen that the simple algorithm assigning each job to its fastest machine is an $O(m)$ competitive algorithm.

We proceed with the unapproximability result:

THEOREM 2.1. *Any online algorithm for the load balancing of temporary tasks on unrelated machines is $\Omega(m/\log m)$ competitive.*

Proof. Let k be the largest integer power of 2 such that $k \le m/\log m$. Assume by contradiction that there is an online algorithm whose competitive ratio is below $k/2$. We describe a sequence of jobs given by an adversary such that there exists an optimal assignment whose maximum load is at most 1. The sequence ends as soon as there is a machine whose load in the online assignment is at least $k/2$.

The lower bound is composed of $l = \log k$ sets of k machines each. The sets are denoted by $M_1, M_2, ..., M_l$. In addition, a machine denoted by m_0^1 is used. Also, denote by m_i^j the j'th machine in the set M_i, $1 \le i \le l$, $1 \le j \le k$. Note that the total number of machines used, $l \cdot k + 1$, does not exceed m (when $m > 2$).

The adversary proceeds in phases. Before the start of phase t, we define a set of $l+1$ machines which we call active. One machine in each M_i is active and we denote its index by $a_i(t)$. The machine m_0^1 is always active and we use the notation $a_0(t) = 1$. The load of $m_i^{a_i(t)}$, $i = 0, ..., l$, in the online assignment is denoted by $b_i(t)$. We begin with setting $a_i(0) = 1$ for $i = 1, ..., l$.

A phase is composed of an arrival of one job and the departure of a set of jobs. The job presented by the adversary has infinite weight on non-active machines. Its weight on $m_i^{a_i(t)}$ is $2^i/k$ for $i = 0, ..., l$. Assume the online algorithm assigns it to machine $m_i^{a_i(t)}$. In case the new load, $b_i(t+1)$, is $k/2$ or more the sequence stops. Otherwise, the phase is completed with the departure of the jobs assigned by the online algorithm to $m_{i+1}^{a_{i+1}}, ..., m_l^{a_l}$ (no jobs leave when $i = l$). The set of active machines for the next phase is set as follows: $a_j(t+1) = 1$ for any $i+2 \le j \le l$ and unless $i = l$ we also set $a_{i+1}(t+1) = 1 + \lfloor 2b_i(t) \rfloor$. All other active machines stay the same. Note that since $b_i(t) < k/2$ the above definition of $a_i(t+1)$ is valid, that is, $a_i(t+1) \le k$. Also note that by the above construction, non-active machines are always empty.

If we consider the vector of loads of the online assignment, $(b_0(t), b_1(t), ..., b_l(t))$ we note that the vector increases lexicographically after each phase. That is, at least one of the coordinates increases while all previous coordinates do not change. The increase is by at least $1/k$. Since the adversary sequence is completed once one of the coordinates exceeds $k/2$, the sequence is completed after a finite number of phases, or specifically, at most $O(k^{2l}) = O(m^{2\log m})$ phases.

In what follows we complete the proof by showing an optimal assignment where the maximum load does not exceed 1 during the whole sequence. In case the job arriving at phase t is assigned by the online algorithm to $m_i^{a_i(t)}$, $i = 0, ..., l-1$, then the optimal algorithm assigns it to machine $m_{i+1}^{a_{i+1}(t)}$. Otherwise, the online algorithm assigns the job to $m_l^{a_l(t)}$ in which case the optimal algorithm assigns the job to machine m_0^1.

First, jobs assigned by the optimal assignment to m_0^1 are assigned by the online algorithm to an active machine in M_l. Since that machine's load is not more than $k/2$ and all other machines in M_l are empty, the incurred load on m_0^1 is at most $1/2$. Now consider jobs assigned to m_{i+1}^j, $i = 0, ..., l-1$, by the optimal assignment. These are assigned by the online algorithm to the active machine in M_i. Moreover, when they were assigned, the load on the active machine in M_i was at least $(j-1)/2$ and less than $j/2$. Therefore their total load on a machine in M_i is at most $1/2$ and their total load on a machine in M_{i+1} is at most 1.

Concluding, the above sequence was shown to have an optimal assignment of maximum load 1. Moreover, as long as the online maximum load is below $k/2$ the online load vector was shown to increase lexicographically. This contradicts our assumption that an online algorithm with competitive ratio below $k/2$ exists and completes the proof.

The following lemma demonstrates a general technique for converting deterministic lower bounds into randomized ones.

LEMMA 2.1. *Let c be a lower bound on the competitive ratio of any deterministic algorithm for unrelated assignment of temporary tasks. If c can be proven with an adversarial strategy of jobs in which each job has an admissible set of at most k machines, then there is a lower bound of $\frac{c}{k}$ on the competitive ratio of any randomized on-line algorithm for the same problem.*

Proof. Omitted.

THEOREM 2.2. *Any online randomized algorithm for the load balancing of temporary tasks on unrelated machines is $\Omega(m/(\log m)^2)$ competitive.*

Proof. The construction in Theorem 2.1 uses admissible sets of at most $\log m$ machines. The theorem then follows as a corollary of Lemma 2.1.

2.2 Tight Results for the Related-Restricted Model

The result in the previous section shows that approximating the unrelated machines model is almost infeasible. As an alternative to the unrelated machines model we consider the so called related-restricted model. Here, each machine has its own speed and each job has a weight and a set of admissible machines. However, note that the lower bound presented in the last section still applies here so approximating is still infeasible. We show that by limiting the number of different machine speeds to a constant number, we can approximate the problem better. Specifically, in the case where only two different machine speeds are involved we obtain the following tight result. More details will appear in the final version of the paper.

THEOREM 2.3. *Any online algorithm for the load balancing of temporary tasks in the related-restricted model with speeds* $\{1, s\}$ *is* $\Omega(\min\{\max\{m/s, \sqrt{m}\}, \sqrt{ms}\})$ *competitive. In addition, there exists an* $O(\min\{\max\{m/s, \sqrt{m}\}, \sqrt{ms}\})$*-competitive algorithm. The same holds for randomized algorithms as well.*

Proof. Omitted.

2.3 Tight Results for Restricted Assignment

In this section we show that the greedy algorithm which assigns a task to the least loaded admissible machine (breaking ties arbitrarily) is optimal for restricted assignment of temporary tasks with unknown durations when we have a small number of machines. This is in contrast to its non-optimal performance for a general m. We first perform an analysis of the greedy algorithm, which gives better upper bounds for small m. It can be seen that greedy is optimal for the case of two machines and achieves a competitive ratio of 2. We show that it is also optimal for three, four and five machines by proving matching lower bounds.

THEOREM 2.4. *The greedy algorithm is optimal for the assignment of temporary tasks in the restricted assignment model for instances where* $m \leq 5$.

Proof. Omitted.

3 Tasks of Known Duration

THEOREM 3.1. *Any deterministic on-line algorithm for load-balancing of temporary tasks with known durations in the restricted assignment model has a competitive ratio of at least* $\sqrt{m}$.

Proof. Let ON be an on-line algorithm for the problem, and let OFF be an optimal offline algorithm for solving it. We will show a sequence of jobs for which ON would reach a load of at least $\sqrt{m}$, where OFF maintains a load of one. First we describe the sequence, and then we prove the lower bound.

We denote the set of the first $\sqrt{m}$ machines by A, and the set of the remaining machines by B. The i'th machine in A (respectively B) will be denoted by A_i (respectively, B_i), for $1 \leq i \leq \sqrt{m}$ (respectively, for $1 \leq i \leq m - \sqrt{m}$).

We force ON to assign $\sqrt{m}$ jobs to a single machine in B, or to assign m jobs to A (which consists of only $\sqrt{m}$ machines).

Our input sequence only includes unit jobs and consists of at most $m - \sqrt{m}$ phases. Each phase p ($p \geq 1$) consists of at most $\sqrt{m}$ jobs. The j'th job in phase p ($j \geq 1$) is admissible to two machines: A_j and B_p. The exact arrival and departure time of each job will be described later. The number of jobs arriving in each phase is determined by the behavior of ON. As long as ON assigns the jobs in phase p to B_p, jobs keep arriving at that phase (up to the maximum of $\sqrt{m}$ jobs). When ON assigns a job to a machine in A, the phase ends (i.e. no more jobs arrive in this phase). Let N_p be the number of jobs which arrived in phase p. By definition, $1 \leq N_p \leq \sqrt{m}$. The number of phases is also determined by the behavior of ON. If $N_p = \sqrt{m}$ for a certain phase p (i.e., ON assigns all the jobs of that phase to B_p), then the sequence stops. If phase $m - \sqrt{m}$ had less than $\sqrt{m}$ jobs, then we bring one more unit job ("extra job"), which will be restricted to the most loaded machine that ON has in A.

We now describe the arrival and departure times of the jobs in each phase. We first describe these times for the first phase, and then we inductively define them for the other phases. The length of the time interval that our sequence will use is $T = \sqrt{m}^{(m-\sqrt{m}+1)}$. Let $S_1 = 0$ and let $T_1 = T$. The first phase starts at time $S_1 = 0$. The jth job of phase 1 arrives at time $j - 1$ ($1 \leq j \leq N_1$). Its departure time is $\frac{j \cdot T_1}{\sqrt{m}}$.

For each phase $p > 1$, we inductively define the arrival and departure times of the jobs to be between the departure times of the last two jobs of the previous phase. For $p \geq 1$, we define T_{p+1} as the departure time of the last (i.e. N_p'th) job of phase p. We also define S_{p+1} as the departure time of the $N_p - 1$'st job of phase p. If $N_p = 1$ then S_{p+1} is equal to S_p. Each phase p starts at time S_p, and only uses the time interval $[S_p, T_p]$. The arrival time of the j'th job in phase p is $S_p + j - 1$, and its departure time is $S_p + \frac{j \cdot (T_p - S_p)}{\sqrt{m}}$. Recall that in case $N_{m-\sqrt{m}} < \sqrt{m}$ we add one more unit job, restricted to the most loaded machine that ON has in A. This "extra job" lasts just one time unit and arrives at time $S_{m-\sqrt{m}+1}$. This completes the description of our sequence.

We first prove that ON achieves a load of at least $\sqrt{m}$ for the above sequence. We notice that the first departure of a job in a certain phase in our sequence occurs only after the arrival of the last job in that phase. We will briefly explain this. The minimal duration of a job in a phase (i.e. the duration of the first job) is divided by $\sqrt{m}$ in each phase. Since the minimal duration in the first phase is $\sqrt{m}^{(m-\sqrt{m})}$, the minimal duration in the last phase (which is the minimal duration of any job during all our phases) is $\sqrt{m}$. Since all the jobs of a certain phase arrive until the $(\sqrt{m}-1)$ moment of that phase, the first departure always occurs after the last arrival.

Therefore, when a job arrives all the previous jobs of its phase are still active. This means that if ON assigns all the jobs in phase p to machine B_p, then it reaches a load of at least $\sqrt{m}$ (all these jobs are active together at time $S_p+\sqrt{m}-1$), and we are done.

In order to avoid this, ON must assign a job to A at a certain stage of each phase. By definition, the phase ends as soon as this happens. Recall that phase $p+1$ starts when the N_p-1'st job of phase p leaves, and ends before the departure of the N_p'th job of phase p. So the last job which arrived in phase p is the only active job from phase p at the beginning of phase $p+1$, and it remains active throughout that phase (i.e. $T_p \geq T_{p+1}$ and $S_{p+1} \geq S_p$). Inductively, $T_{p-1} \geq T_p$ and $S_p \geq S_{p-1}$. Thus, the last job of each of the phases $1,...,p$ is still active throughout phase $p+1$, and these last jobs are the only jobs from phases $1,...,p$ which are active at phase $p+1$. Recall that the last job in each phase is the job that ON assigned to A. So at the beginning of phase p (time S_p), the active jobs are exactly all the jobs that ON assigned to A in phases $1...(p-1)$. Summing over all the phases, on time $S_{m-\sqrt{m}+1}$, ON has $m-\sqrt{m}$ jobs in A. There are only $\sqrt{m}$ machines in this set, so the most loaded machine in A, machine A_i, has a load of at least $\sqrt{m}-1$. As we explained before, the "extra job" now arrives, and can only be assigned to machine A_i. This makes the load of ON on that machine at least $\sqrt{m}$.

Now we will describe the assignment of algorithm OFF. The strategy of OFF is simple. When the j'th job of phase p arrives, it can be assigned either to a machine from A, A_j, or to a machine from B, B_p. If ON assigns the job to the machine in B, then OFF assigns it to the machine in A. If ON assigns it to the machine in A, then OFF assigns it to the machine in B. If the "extra job" arrives, then OFF assigns it to its admissible machine.

Let us consider now the load of OFF. At the beginning of phase p, OFF has no active jobs in A, since we saw that ON has no active jobs in B. OFF has one active job on each of the machines $B_1,...,B_{p-1}$, since ON has one active job from each phase in A. During phase p, as long as ON assigns jobs to B_p, OFF assigns each of them to a different machine in A (which was empty at the beginning of the phase). When ON assigns a job to A, OFF assigns it to B_p (which is empty), and the phase ends (so no other job will be assigned to B_p). Therefore, OFF maintains a load of 1 throughout the phases. At time $S_{m-\sqrt{m}+1}$, OFF has one active job on each machine in B (one job from each phase), and no jobs in A. So it can assign the "extra job" to A_i without exceeding the maximum load of 1. Thus we have reached the required competitive ratio.

Let us denote the total length of the time interval used by the input sequence by T. The result above also applies when we limit the length T of the sequence and when we allow randomization to be used. The results are summarized in the next two theorems:

THEOREM 3.2. *Any deterministic on-line algorithm for load-balancing of temporary tasks with known durations in the restricted assignment model has a competitive ratio of at least* $\Omega(\sqrt{\frac{\log T}{\log\log T}})$, *for any* $T < \sqrt{m}^{(m-\sqrt{m}+1)}$.

Note that this lower bound is at most $\sqrt{m}$ for this range of T.

Proof. For any $T = \sqrt{x}^{x-\sqrt{x}+1}$ where $x \leq m$, we can clearly apply the exact steps of the previous proof, limiting ourselves to the first x machines instead of using all the machines. We will have a sequence with at most $x-\sqrt{x}$ phases, each of them having at most $\sqrt{x}$ jobs, and we will obtain a lower bound of $\sqrt{x}$. We can easily see that in this case: $\sqrt{x} = \Omega(\sqrt{\frac{\log T}{\log\log T}})$. This means that for any $T < \sqrt{m}^{(m-\sqrt{m}+1)})$, we have a lower bound of $\Omega(\sqrt{\frac{\log T}{\log\log T}})$, as required.

THEOREM 3.3. *A randomized on-line algorithm for solving the problem of restricted assignment of temporary tasks with known durations cannot achieve a competitive ratio smaller than* $\frac{1}{2}\sqrt{m}$. *Moreover, for any* $T < \sqrt{m}^{(m-\sqrt{m}+1)}$ *no algorithm can be better than* $\Omega(\sqrt{\frac{\log T}{\log\log T}})$*-competitive.*

Proof. Note that in Theorems 3.1 and 3.2 admissible sets contain at most two machines. The results follow by using Lemma 2.1.

4 Off-line Temporary Assignment

4.1 Fixed Number of Machines

We briefly describe the polynomial-time approximation scheme and leave the details to the appendix. We begin with scaling the weights of the jobs, in order to limit the possible range of the optimal maximum load. It is well-known that we can achieve an approximation ratio of m simply by assigning each job to its fastest machine. We will refer to this simple algorithm as "Fastest-Assign". We apply this algorithm to our input, and denote the maximum load reached by l. Now we multiply each of our jobs' weights by $\frac{m}{l}$. This assures us that the optimal maximum load is in the range $[1, m]$. Note that this scaling requires only linear time.

The algorithm then follows with five phases: the weight-rounding and grouping phase, the time- rounding phase, the combining phase, the solving phase and the converting phase. In the first phase, the weights of the jobs are rounded upwards, and then they are divided into a large number of subsets based on their rounded weights, as will be explained later. Next the time-rounding phase is applied to each of these subsets. This phase actually consists of two subphases. In the first subphase the jobs' active time is extended: some jobs will arrive earlier, others will depart later. In the second subphase, the active time is again extended but each job is extended in the opposite direction to which it was extended in the first subphase. The combining phase is also applied to each subset separately. In this phase the algorithm combines several jobs from the same subset into jobs with higher load vector coordinates. In the solving phase, we find an optimal solution for the modified problem (the solving is performed for all the jobs together). The solution we found can be converted into a solution for the original problem in the converting phase, which is again applied separately to each subset.

THEOREM 4.1. *The algorithm described in Appendix A.1 is a PTAS running in time* $O(n^{1+\epsilon^{-6}m^7(\lceil\log_{1+\epsilon}(m/\epsilon)\rceil+1)^m \log m})$.

Proof. Omitted.

4.2 Non-Fixed Number of Machines

In this section we consider offline load-balancing of temporary tasks on a non-fixed number of unrelated machines. We show that no polynomial approximation algorithm can achieve an approximation ratio smaller than 2, unless $P = NP$. We prove this lower bound for the special case of restricted assignment (so it obviously holds for the general case of any unrelated machines as well).

THEOREM 4.2. *For every $\rho < 2$, there does not exist a polynomial ρ-approximation algorithm for restricted assignment of temporary tasks unless $P = NP$.*

Proof. We use a reduction from the 3-dimensional-matching problem ($3DM$), which is known to be NP-complete. In that problem, we are given three sets of elements, B, G, and H, each of them of size n ($B = b_1, ..., b_n$, $G = g_1, ..., g_n$, $H = h_1, ..., h_n$). We are also given a set $S = T_1, ..., T_m$ of m triplets, $S \subseteq B * G * H$. These are the possible matchings of 3 elements from B,G, and H. The goal is to decide whether there exists a matching for all the elements of B,G, and H, i.e. a subset of S, S', such that $|S'| = n$ and $\bigcup_{T_i \in S'} T_i = B \cup G \cup H$. Given an instance to the $3DM$ problem we construct an instance for our problem. For each triplet T_i, we have a machine m_i. All our jobs are of weight 1, and our total time interval will be of length 3 (from time 0 until time 3). Our first type of jobs will be "element jobs" (one job for each element), as described hereafter. For each element $b_i \in B$, we will have a job which arrives at time 0, departs at time 1, and is admissible to a machine m_k if and only if $b_i \in T_k$. For each element $g_i \in G$, we will have a job which arrives at time 1, departs at time 2, and is admissible to a machine m_k if and only if $g_i \in T_k$. Finally, for each element $h_i \in H$, we will have a job which arrives at time 2, departs at time 3, and is admissible to a machine m_k if and only if $h_i \in T_k$. Our second type of jobs will constitute of $m - n$ "dummy jobs", which arrive at time 0, depart at time 3, and are admissible to all the machines.

We prove that there is an assignment with a maximum load of 1, if and only if there is a solution to the $3DM$ problem. Suppose there is a $3DM$, $S' = \{T_{i_1}, ..., T_{i_n}\}$. Then for each $T_{i_k} \in S'$, we assign to the machine m_{i_k} the 3 "element jobs" which correspond to the 3 elements of T_{i_k}. They are admissible to m_{i_k}, because this is how we defined our assignment restrictions. Also notice that they are active in different times, so the machine maintains a load of 1. This way we assign the $3n$ "element jobs" to n of the machines. We assign the $m - n$ "dummy jobs" to the other $m - n$ machines, one job on each machine. They are admissible on any machine, and have a weight of 1 each. Therefore, this assignment maintains a maximum load of 1 as required.

Now, assume that there is an assignment having a maximum load of 1. The $m - n$ "dummy jobs" must have been assigned to $m - n$ different machines (if there is more than one "dummy job" on a machine, then its load is bigger than 1). A "dummy job" is active during our entire time interval, so a machine which has a "dummy job" on it cannot have any other job assigned to it. Therefore, the $3n$ "element jobs" must have been assigned to the remaining n machines. Each of these machines $m_{i_1}, ..., m_{i_n}$ must have one active "element job" on it at each moment (since the total volume of the "element jobs" is $3n$). This means that each of these machines, m_{i_k}, has a job which corresponds to an element of B assigned to it, a job which corresponds to an element of G assigned to it, and a job which corresponds to an element of H assigned to it (this is the only possibility to have an active "element job" at each moment). According to our assignment restrictions, these 3 elements from B,G, and H must be included in the triplet T_{i_k}. All of the "element jobs" were assigned, so $\bigcup_{T_{i_k}} T_i = B \cup G \cup H$, and therefore $T_{i_1}, ..., T_{i_n}$ is a $3DM$.

The above reduction shows that any approximation algorithm for our problem with an approximation ratio strictly less than 2 solves the $3DM$ problem. Hence, we have proven the theorem.

References

[1] S. Albers. Better bounds for on-line scheduling. In *Proc. 29th ACM Symp. on Theory of Computing*, pages 130–139, 1997.

[2] Jim Aspnes, Yossi Azar, Amos Fiat, Serge Plotkin, and Orli Waarts. On-line load balancing with applications to machine scheduling and virtual circuit routing. In *Proc. 25th ACM Symp. on Theory of Computing*, pages 623–631, May 1993.

[3] Y. Azar. On-line load balancing. In A. Fiat and G. Woeginger, editors, *Online Algorithms - The State of the Art*, chapter 8, pages 178–195. Springer, 1998.

[4] Y. Azar, A. Z. Broder, and A. R. Karlin. On-line load balancing. *Theoretical Computer Science*, 130(1):73–84, 1994. Also in *Proc. 33rd IEEE FOCS*, 1992, pp. 218-225.

[5] Y. Azar and L. Epstein. On-line load balancing of temporary tasks on identical machines. In *5th Israeli Symp. on Theory of Computing and Systems*, pages 119–125, 1997.

[6] Y. Azar, B. Kalyanasundaram, S. Plotkin, Kirk R. Pruhs, and Orli Waarts. On-line load balancing of temporary tasks. *Journal of Algorithms*, 22(1):93–110, 1997. Also in *Proc. WADS'93*, pp. 119-130.

[7] Y. Azar, J. Naor, and R. Rom. The competitiveness of on-line assignments. *Journal of Algorithms*, 18(2):221–237, 1995. Also in *Proc. 3rd ACM-SIAM SODA*, 1992, pp. 203-210.

[8] Y. Azar, O. Regev, J. Sgall, and G. Woeginger. Off-line temporary tasks assignment. *Theoretical Computer Science*. To appear. Also in *Proc. 7th Annual European Symposium on Algorithms* 1999, pp. 163-171.

[9] B. S. Baker, D. J. Brown, and H. P. Katseff. A 5/4 algorithm for two-dimensional packing. *J. Algorithms*, 2:348–368, 1981.

[10] Y. Bartal, A. Fiat, H. Karloff, and R. Vohra. New algorithms for an ancient scheduling problem. In *Proc. 24th ACM Symposium on Theory of Algorithms*, pages 51–58, 1992. Also in *Journal of Computer and System Sciences* (1995) 359-366.

[11] P. Berman, M. Charikar, and M. Karpinski. A note on on-line load balancing for related machines. In *5th annual Workshop on Algorithms and Data Structures*, 1997.

[12] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.

[13] B. Chen, A. van Vliet, and G. J. Woeginger. New lower and upper bounds for on-line scheduling. *Operations Research Letters*, 16:221–230, 1994.

[14] W. Fernandez de la Vega and G. S. Lueker. Bin packing can be solved within $1 + \epsilon$ in linear time. *Combinatorica*, 1:349–355, 1981.

[15] M. R. Garey, R. L. Graham, D. S. Johnson, and A. C. C. Yao. Resource constrained scheduling as generalized bin packing. *J. Comb. Th. Ser. A.*, 21:257–298, 1976.

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

[17] E. Horowitz and S. Sahni. Exact and approximate algorithms for scheduling non-identical processors. *Journal of the Association for Computing Machinery*, 23:317–327, 1976.

[18] D. Karger, S. Phillips, and E. Torng. A better algorithm for an ancient scheduling problem. In *Proc. 5th ACM-SIAM Symp. on Discrete Algorithms*, pages 132–140, 1994.

[19] J.K. Lenstra, D.B. Shmoys, and E. Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Math. Prog.*, 46:259–271, 1990.

[20] S. Phillips and J. Westbrook. On-line load balancing and network flow. In *Proc. 25th ACM Symposium on Theory of Computing*, pages 402–411, 1993.

[21] S. Plotkin and Y. Ma. An improved lower bound for load balancing of tasks with unknown duration. *Inform. Process. Lett.*, 62:301–303, 1997.

[22] D. Shmoys and E. Tardos. An approximation algorithm for the generalized assignment problem. *Mathematical Programming A*, 62:461–474, 1993. Also in the proceeding of the 4th Annual ACM-SIAM Symposium on Discrete Algorithms, 1993.

[23] J. Westbrook. Load balancing for response time. In *3rd Annual European Symposium on Algorithms*, 1995.

A Appendix

A.1 A $PTAS$ for temporary assignment of unrelated tasks

We denote the sequence of events by $\sigma = \sigma_1, ..., \sigma_{2n}$, where each event is an arrival or a departure of a job; we assume that at each time only one job arrives or departs. Since all the events are known at the beginning, we view

σ as a sequence of times, the time σ_i is the moment after the i'th event happened. In addition, σ_0 denotes the moment at the beginning, before the arrival of the first job. We assume without loss of generality that $m \geq 2$ (otherwise the approximation ratio is always 1).

Let $\epsilon' > 0$ be the precision required by the PTAS. We assume that $\epsilon' < 1$. We choose $\epsilon = \epsilon'/7$, and fix the following 3 constants:

$$\alpha = \frac{\epsilon^2}{m\lceil \log n \rceil \cdot (\lceil \log_{1+\epsilon}(\frac{m}{\epsilon})\rceil + 1)^m}$$

$$\beta = \frac{\alpha\epsilon^2}{m^2} = \frac{\epsilon^4}{m^3\lceil \log n \rceil \cdot (\lceil \log_{1+\epsilon}(\frac{m}{\epsilon})\rceil + 1)^m}$$

$$\gamma = \frac{\beta\epsilon^2}{m^2} = \frac{\epsilon^6}{m^5\lceil \log n \rceil \cdot (\lceil \log_{1+\epsilon}(\frac{m}{\epsilon})\rceil + 1)^m}$$

Phase 1: The weight-rounding and grouping phase. We start by describing the weight-rounding and grouping phase. For each job j, we denote by W_j the load that j causes on its fastest machine: $W_j = min_i(\bar{w}_j(i))$. We will refer to W_j as the "min-weight" of the job j. We define the "relative speed" vector of job j, $\bar{v}_j$ by: $\bar{v}_j(i) = \bar{w}_j(i)/W_j$, for $1 \leq i \leq m$. Note that $\bar{v}_j(i) \geq 1$. We now perform a rounding of the vector $\bar{v}_j$ and obtain the rounded "relative speed" vector $\bar{v}'_j$. For each $1 \leq i \leq m$, if $\bar{v}_j(i) \geq m/\epsilon$, then $\bar{v}'_j(i) = \infty$, and we will refer to machine i as an "illegal machine" for job j. Otherwise, we obtain $\bar{v}'_j(i)$ by rounding $\bar{v}_j(i)$ upwards to the nearest power of $1+\epsilon$, and we will refer to machine i as a "legal machine" for job j. Note that each coordinate of the vector $\bar{v}'_j$ may have $\lceil \log_{(1+\epsilon)}(\frac{m}{\epsilon})\rceil + 1$ possible values, since its value is either ∞ or a power of $(1+\epsilon)$ between 1 and m/ϵ. Now we define a new loads vector for the job j, $\bar{w}'_j$, by: $\bar{w}'_j(i) = W_j \cdot \bar{v}'_j(i)$, for $1 \leq i \leq m$. In this we completed the rounding of the weights.

Next we divide the jobs into subsets according to their $\bar{v}'$ vector. This division splits the jobs into at most $(\lceil \log_{1+\epsilon}(\frac{m}{\epsilon})\rceil + 1)^m$ subsets, since each of the m coordinates of that vector may have $(\lceil \log_{1+\epsilon}(\frac{m}{\epsilon})\rceil + 1)$ possible values, as we noted before. This completes the description of the first phase.

Phase 2: The time-rounding phase. This phase is similar to the time-rounding phase described by [8]. We will apply this phase separately to each subset of jobs $J_{\bar{v}'}$ (having the same "relative speeds" vector $\bar{v}'$).

In order to describe the time-rounding phase with its two subphases, we start with defining partitions of each subset $J_{\bar{v}'}$, based on which the rounding will be performed. The set $R_{\bar{v}'}$ contains all jobs with $W_j \geq \gamma$ out of the jobs in $J_{\bar{v}'}$.

From now on, we fix $\bar{v}'$ in the description of this phase, and refer to $J_{\bar{v}'}$ as J and to $R_{\bar{v}'}$ as R. All the following definitions are made for these fixed J and R.

We begin by defining a partition $\{J_i\}$ of the set of jobs $J - R$. We set $M_1 = J - R$ and define sets J_i and M_i iteratively as follows. Let M_i be a set of jobs and consider the sequence of times in σ in which jobs of M_i arrive and depart. The number of such times is $2r$ for some r, let c_i be any time between the r'th and the $r+1$-st elements in that set. The set J_i contains the jobs in M_i that are active at time c_i. The set M_{2i} contains the jobs in M_i that depart before or at c_i and the set M_{2i+1} contains the jobs in M_i that arrive after c_i. We stop when all unprocessed M_i's are empty. The important property of that partition is that the set of jobs from $J - R$ that are active at a certain time is partitioned into at most $\lceil \log n \rceil$ different sets J_i.

We continue by further partitioning the set J_i. We order the jobs according to their arrival time. We denote the smallest prefix of the jobs whose total min-weight is at least α by S_i^1. We order the same jobs according to their departure time. We take the smallest suffix whose min-weight is at least α and denote that set by T_i^1. Note that there might be jobs that are both in S_i^1 and T_i^1. We remove the jobs in $S_i^1 \cup T_i^1$ from J_i, repeat the process with the jobs left in J_i and similarly define $S_i^2, T_i^2, \ldots, S_i^{k_i}, T_i^{k_i}$. Each set S_i and T_i has total min-weight between α and $\alpha + \gamma$, except for the last pair which may have smaller min-weight than α. However, if the last pair has smaller min-weight than α, then it satisfies $S_i^{k_i} = T_i^{k_i}$. We denote by s_i^j the arrival time of the first job in S_i^j and by t_i^j the departure time of the last job in T_i^j. Note that $s_i^1 \leq s_i^2 \leq \ldots \leq s_i^{k_i} \leq c_i \leq t_i^{k_i} \leq \ldots \leq t_i^2 \leq t_i^1$.

The first subphase of the time-rounding phase creates a new set of jobs J' which contains the same jobs as in J with slightly longer active times. We change the arrival time of all the jobs in S_i^j for $j = 1, \ldots, k_i$ to s_i^j. Also, we change the departure time of all the jobs in T_i^j to t_i^j. The jobs in R are left unchanged. We denote the sets resulting from the first subphase by J', J'_i, S'^j_i, T'^j_i.

The second subphase of the time-rounding phase further extends the active time of the jobs resulting from the first subphase. We take one of the sets J'_i and the partition we defined earlier to $S'^1_i \cup T'^1_i$, $S'^2_i \cup T'^2_i$, $\ldots$, $S'^{k_i}_i \cup T'^{k_i}_i$. For every $j \leq k_i$, we order the jobs in S'^j_i according to an increasing order of departure times. We take the smallest prefix of this ordering whose total min-weight is at least β. We extend the departure time of all the jobs in that prefix to the departure time of the last job in that prefix. The process is repeated until there are no more jobs in S'^j_i. The last prefix may have a min-weight of less than β. Similarly, we extend the arrival times of jobs in ${T'}_i^j$. Note that if the total min-weight of either ${S'}_i^{k_i}$ or ${T'}_i^{k_i}$ is smaller than α then ${S'}_i^{k_i} = {T'}_i^{k_i}$ and these jobs are left unchanged since they already have identical arrival and departure times from the first phase. We denote the sets resulting from the second subphase by J'', J_i'', ${S''}_i^j$, ${T''}_i^j$.

Phase 3: The combining phase. This phase involves the load vectors of the jobs. It is also applied to each subset $J_{\bar{v}'}''$ separately, so we again fix $\bar{v}'$ in the description of this phase and refer to $J_{\bar{v}'}''$ as J''. Let J_{st}'' be the set of jobs in J'' that arrive at s and depart at t. Assume the total min-weight of jobs in J_{st}'' is x. The combining phase replaces these jobs by $\lceil x/\gamma \rceil$ jobs, which have a load-vector of $\gamma \cdot \bar{v}'$. Note that the maximum finite weight in their loads vector may be $\frac{m}{\epsilon} \cdot \gamma$. We denote the resulting sets by J_{st}'''. The set J''' is created by replacing every J_{st}'' with its corresponding J_{st}''', that is, $J''' = \bigcup_{s,t} J_{st}'''$.

Phase 4: The solving phase. This phase solves the modified decision problem, i.e. it solves the problem after each subset $J_{\bar{v}'}$ has been replaced by a modified subset $J_{\bar{v}'}'''$. The solving phase is performed once for all the jobs together (not for each $J_{\bar{v}'}'''$ subset separately). We solve the modified decision problem by building a layered graph. Every time σ_i, $i = 0, \ldots, 2n$, in which jobs arrive or depart (including the initial state with no job) has its own set of vertices called a layer. Each layer holds a vertex for every possible assignment of the current active jobs to machines (except assignments of weight ∞); furthermore, we label each node by the maximum load of a machine in that configuration.

Two vertices of adjacent layers σ_{i-1} and σ_i, $i = 1, \ldots, 2n$, are connected by an edge if the transition from one assignment of the active jobs to the other is consistent with the arrival and departure of jobs at time σ_i. More precisely, the vertices are connected if and only if every job active both before and after σ_i is assigned to the same machine in the assignments of both vertices. At each event, jobs either arrive or depart but not both (due to the assumption at the beginning that all the original events are distinct; during rounding we do not mix arrival and departure events). If σ_i is an arrival, the indegree of all vertices in the layer σ_i is 1, since the new configuration determines the old one. Similarly if σ_i is a departure, the outdegree of all vertices in the layer σ_{i-1} is 1. In both cases, the number of edges between two layers is linear in the number of vertices on these layers. It follows that the total number of edges is linear in the number of vertices.

We define a value of a path as the maximal value of its nodes. Now we can simply find a path with smallest value from the first layer to the last one by any shortest path algorithm in linear time (since the graph is layered).

Phase 5: The converting phase. In this phase the algorithm converts the assignment found for the modified problem into an assignment for the original problem. This phase is performed separately for the jobs in each of the subsets $J_{\bar{v}'}'''$. Each assignment of the jobs of a modified subset $J_{\bar{v}'}'''$ is converted into an assignment for the jobs of subset $J_{\bar{v}'}$, which is also an assignment for the original problem. Again we fix $\bar{v}'$ throughout the description of this phase, and refer to $J_{\bar{v}'}'''$ as J'''. Assume the number of jobs having $W_j = \gamma$ in J_{st}''' that are assigned to a certain machine i is r_i. Remove these jobs and assign all the jobs having $W_j \leq \gamma$ in J_{st}'' to the machines such that a total weight of at most $(r_i + 1)\gamma \cdot \bar{v}'(i)$ is assigned to machine i.

Note that all the jobs will be assigned that way. The replacement involves jobs whose min-weight is at most γ. We know that the total min-weight of these jobs is at most $\gamma \cdot \sum_{i=1}^m r_i$.

If they made a load of $(r_i + 1)\bar{v}'(i)\gamma$ on each of the machines, then it would mean that their total min-weight was at least $\gamma \cdot (m + \sum_{i=1}^m r_i)$. So it is possible to assign all these jobs so that they will make a load of at most $(r_i + 1)\gamma \cdot \bar{v}'(i)$ on each machine i.

The assignment for J'' is also an assignment for J' and J. An assignment for J is also an assignment for the original problem.