Minimizing Maximum Flow-time on Related Machines*

Nikhil Bansal and Bouke Cloostermans

Eindhoven University of Technology
P.O. Box 513, 5600 Eindhoven, The Netherlands
{b.cloostermans,n.bansal}@tue.nl

Abstract

We consider the online problem of minimizing the maximum flow-time on related machines. This is a natural generalization of the extensively studied makespan minimization problem to the setting where jobs arrive over time. Interestingly, natural algorithms such as Greedy or Slow-fit that work for the simpler identical machines case or for makespan minimization on related machines, are not $O(1)$-competitive. Our main result is a new $O(1)$-competitive algorithm for the problem. Previously, $O(1)$-competitive algorithms were known only with resource augmentation, and in fact no $O(1)$ approximation was known even in the offline case.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Related machines scheduling, Maximum flow-time minimization, On-line algorithm, Approximation algorithm

Digital Object Identifier 10.4230/LIPIcs.APPROX-RANDOM.2015.85

1 Introduction

Scheduling a set of jobs on machines to optimize some quality of service measure is one of the most well studied problems in computer science. A very natural measure of service received by a job is the flow-time, defined as the amount of time the job spends in the system. In particular, if a job j arriving at time r_j completes its processing at time C_j, then its flow-time F_j is defined as $C_j - r_j$; i.e., its completion time minus its arrival time. Over the last few years, several variants of flow-time related problems have received a lot of attention: on single and multiple machines, in online or offline setting, for different objectives such as total flow-time, ℓ_p norms of flow-time, stretch etc., with or without resource augmentation, in weighted or unweighted setting and so on. We refer the reader to [15, 13, 12, 4] for a survey of some of these results.

In this paper we focus on the objective of minimizing the maximum flow-time. This is desirable when we want to guarantee that *each* job has a small delay. Maximum flow-time is also a very natural generalization of the minimum makespan or the load-balancing problem, that has been studied extensively (see e.g. [5, 9, 1] for a survey). In particular, if all jobs have identical release times, then the maximum flow-time value is precisely equal to the makespan. Minimizing the maximum flow-time is also related to deadline scheduling problems. In particular, the maximum flow-time is at most D if and only if each job j completes by $r_j + D$. Moreover, note that arbitrary deadlines d_j can be modeled by considering the weighted version of maximum flow-time[1].

* Supported by NWO grant 639.022.211 and ERC consolidator grant 617951.

[1] In deadline scheduling however, the deadlines are typically considered fixed and the focus is on maximizing the throughput.

18th Int'l Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'15) / 19th Int'l Workshop on Randomization and Computation (RANDOM'15).
Editors: Naveen Garg, Klaus Jansen, Anup Rao, and José D. P. Rolim; pp. 85–95
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Known results for maximum flow-time

For a single machine, it is easy to see that First In First Out (FIFO) is an optimal (online) algorithm for minimizing the maximum flow-time. For identical multiple machines, Bender et al. [8] showed that the Greedy algorithm, that schedules the incoming job on the least loaded machine is $3 - 2/m$ competitive, where m is the number of machines. They also showed that this bound is tight for the Greedy algorithm. If jobs can be preempted and migrated (moved from one machine to another), [2] gave a 2-competitive algorithm.

A systematic investigation of the problem for various machine models was initiated recently by Anand et al. [3]. Recall that in the related machines model each machine i has speed s_i, and processing job j on machine i takes $p_{ij} = p_j/s_i$ units of time. In the more general unrelated machines model, p_{ij} can be completely arbitrary.

Among other results, Anand et al. [3] gave a $(1 + \epsilon)$-speed $O(1/\epsilon)$-competitive algorithm for the unrelated machine case, for any given $\epsilon > 0$. Here the online algorithm can process $1 + \epsilon$ units of work per time step, but is compared to an offline optimum that does not have this extra *resource augmentation* [14, 15]. They also showed that in the unrelated setting any algorithm without resource augmentation must be $\Omega(m)$ competitive. For the weighted maximum flow-time objective, they gave a $(1 + \epsilon)$-speed, $O(1/\epsilon^3)$-competitive algorithm for the related machines setting, and showed that no $O(1)$-speed, $O(1)$-competitive algorithm exists in the unrelated setting.

A natural question that remains is the complexity of the problem for related machines: Is there an $O(1)$-competitive algorithm for the related machines setting, without using resource augmentation?

This question is particularly intriguing as it is not at all clear what the right algorithm should be [2]. In fact, no $O(1)$-approximation is known even in the offline case. One issue is that the natural Slow-fit algorithm, that is $O(1)$-competitive for makespan minimization (even when the jobs are temporary and have unknown durations [6]), is not $O(1)$-competitive for maximum flow-time (Lemma 2 below). The algorithm of [3] for weighted maximum flow-time with resource augmentation is also a variant of Slow-fit. Recently, [7] obtained an $O(\log n)$ approximation for minimizing maximum flow-time on unrelated machines, where n is the number of machines. However their techniques do not seem to give anything better for the related machines setting either.

Our main result is the following.

▶ **Theorem 1.** *There is a* 13.5 *competitive algorithm,* Double-fit*, for minimizing maximum flow-time on related machines.*

This also gives the first $O(1)$ approximation for the offline problem. We also show that no such result is possible in the weighted case (without resource augmentation), and give an $\Omega(W)$ lower bound on the competitive ratio where W is the maximum to minimum weight ratio.

High-level approach

There are two competing trade-offs while scheduling on related machines. On one hand the algorithm should keep as many machines busy as possible, otherwise load might accumulate and delay future jobs. This accumulated load could be impossible to get rid of if there is no resource augmentation. On the other hand, the algorithm should keep fast machines empty for processing large jobs that might arrive later. In particular, fast machines are a scarce resource that should not be wasted on processing small jobs unnecessarily. It is instructive to consider the lower bounds in Section 2, where both Slow-fit and Greedy are shown to perform badly due to these opposite reasons.

To get around this, we design an algorithm that combines the good properties of both Slow-fit and Greedy. In particular, the algorithm uses a two phase strategy while assigning jobs to machines at each step. First, the jobs are spread out to ensure that machines are busy as much as possible. Once machines are *saturated*, the algorithm shifts into a *Slow-fit* mode, which ensures that small jobs do not unnecessarily go on fast machines.

The key difficulty in the analysis is to control how the two phases interact with each other. To do this, we maintain two invariants that capture the dynamics of the algorithm, and control how much the online algorithm's load on a subset of machines deviates from the offline algorithm's load on those machines. The main part of the argument is to show inductively that these invariants are maintained over time.

Notation and formal problem description

There are m machines indexed by non-decreasing order of speeds $s_1 \leq s_2 \leq \ldots \leq s_m$. The processing requirement of job j is p_j, and it requires time p_j/s_i on machine i. We will call p_j the work of j, and p_j/s_i its load on machine i. Jobs arrive online over time and p_j is known immediately upon its release time r_j. The goal is to find a schedule that minimizes the maximum flow-time, and we assume that a job cannot be migrated from one machine to another. We use Opt to denote some fixed optimum offline schedule, and also to denote the value of this solution.

2 Lower bounds on Slow-fit and Greedy

Slow-fit

Algorithm Slow-fit takes as input a threshold F_{opt} (the current guess on optimum), and schedules every incoming job on the slowest possible machine while keeping the load below F_{opt}. If the jobs cannot be feasibly scheduled on any machine, the algorithm fails and the threshold is doubled.

▶ **Lemma 2.** Slow-fit *has a competitive ratio of* $\Omega(m)$*.*

Proof. We describe an instance where the threshold F_{opt} keeps doubling until it reaches m even though Opt $= 2$.

There are m identical machines (but we arbitrarily order them from slow to fast). Next, we assume that $F_{\text{opt}} \geq 2$, which can be achieved by giving $2m$ unit-size jobs initially at $t = 0$.

At each time step $t \geq 2$, m unit-length jobs arrive. As Slow-fit will not use all m machines initially, there will be some time t_0 at which all the machines $1, \ldots, m-1$ have load F_{opt}. At time $t_0 + 1$, when these initial $m-1$ machines have $F_{\text{opt}} - 1$ pending jobs, we release $2m$ unit-size jobs. As there is at most $m - 1 + F_{\text{opt}}$ total capacity available, these jobs cannot be scheduled feasibly if $F_{\text{opt}} \leq m$. On the other hand, at each time step Opt distributes the incoming jobs over all machines and achieves value 2. ◀

Intuitively, Slow-fit unnecessarily builds up load on slow machines while keeping the fast machines empty, and cannot recover if there is small burst of jobs.

Greedy

When a job j arrives, Greedy schedules j on the machine that minimizes the flow-time of j (assuming FIFO order). Ties are broken arbitrarily. The following bound is well-known [11], but we sketch it here for completeness. The idea is that Greedy puts too many slow jobs on fast machines, which causes problems when large jobs arrive.

▶ **Lemma 3.** Greedy *has a competitive ratio of* $\Omega(\log m)$.

Proof. Consider an instance where we have k groups of machines where group G_i contains 2^{2k-2i} machines of speed 2^i. Note that the total processing power in group G_i is equal to $S_i = 2^{2k-i}$. The processing power of groups $i, \ldots, k$ combined is thus equal to $P_i = \sum_{i'=i}^{k} 2^{2k-i'} \leq 2S_i$.

We receive k sets of jobs, all at time 0, but in order. For all $i = 1, \ldots, k$, set J_i contains 2^{2k-2i} jobs of size 2^i. Again, note that the total size of jobs in set J_i is equal to 2^{2k-i}. Greedy will spread jobs from set i over groups $i, \ldots, k$. Group k (containing only a single machine of speed 2^k) will receive a $S_k/P_i \geq \frac{1}{2}S_k/S_i = 2^{-k+i-1}$ fraction of these jobs. This means group k receives $\sum_{i=1}^{k} 2^{2k-i}2^{-k+i-1} = k2^{k-1}$ work. Since group k has a single machine of speed 2^k, finishing these jobs takes $\Omega(k)$ time.

However, optimum can schedule the i-th batch of jobs on group i machines, incurring a maximum load of 1 (i.e., it does Slow-fit with threshold 1). ◀

3 The Algorithm Double-fit

We describe our algorithm, denoted by Double-fit hereafter. Double-fit takes an input a parameter F_{opt}, which is supposed to be our estimate of Opt. By a slight variation on the doubling trick that loses an additional factor of 1.5 (see Section 3.4), we will assume henceforth that $F_{\text{opt}} \in [\text{Opt}, 1.5\text{Opt})$.

We divide time into intervals I_k of size $3F_{\text{opt}}$ as $I_k = [3(k-1)F_{\text{opt}}, 3kF_{\text{opt}})$. We refer to time $3kF_{\text{opt}}$ as the k-th epoch. For each $k = 1, 2, \ldots,$ Double-fit batches the jobs that arrive during I_k and schedules them at epoch k using the algorithm in Figure 1. We use $[i : m]$ to denote the machines $i, \ldots, m$. If the total remaining work on jobs on machine i is $w(i)$ at time t, we say that it has load $w(i)/s_i$.

1. Let J denote the set of jobs arriving during I_k.
2. Partition jobs in J into classes $J_1, \ldots, J_m$, where each job j is in class J_i with the smallest index i such that $p_j \leq s_i \cdot F_{\text{opt}}$.
3. **For** $i = m, m-1, \ldots, 1$
4. Consider the jobs j in J_i in arbitrary order and assign them as follows:
5. **(Saturation Phase:) If** some machine in $[i : m]$ is loaded below $3F_{\text{opt}}$
6. schedule j on the slowest such machine.
7. **(Slow-fit Phase:) Else** schedule j on the slowest machine in $[i : m]$
8. such that its load stays below $6F_{\text{opt}}$.
9. **If** no such machine exists return FAIL.

Figure 1 Algorithm Double-fit for the epoch k.

Description

First, Double-fit classifies the jobs arriving during I_k depending on the smallest machine on which they have size no larger than F_{opt}. Note that as $F_{\text{opt}} \geq \text{Opt}$, if job j is put in class J_i, then Opt cannot schedule job j onto a machine smaller than i either.

Double-fit considers jobs from classes J_m down to J_1 (this ordering will be used crucially). Each class is scheduled in two phases. In the saturation phase, when scheduling a job j, it checks if there is some machine in $[i : m]$ with load less than $3F_{\text{opt}}$. If so, j is scheduled on the slowest such machine. If no such machine exists, the algorithm enters the Slow-fit phase (for class J_i), and performs Slow-fit for class J_i on machines $[i : m]$ with threshold $6F_{\text{opt}}$.

3.1 Analysis

Our goal in this section is to show the following result.

▶ **Theorem 4.** *If $F_{opt} \geq Opt$, then the algorithm never fails.*

This directly implies Theorem 1 as follows. Each job spends at most $3F_{\text{opt}}$ time waiting to be assigned, and at most $6F_{\text{opt}}$ on its designated machine, thus the flow-time of any job is at most $9F_{\text{opt}}$. As $F_{\text{opt}} \leq 1.5\text{Opt}$ by the doubling trick, this implies a competitive ratio of 13.5

For the purpose of analysis, it will be convenient to consider a *restricted* Opt that also batches jobs and schedules the jobs arriving in I_k at epoch k. Note that such a restricted algorithm has objective at most $3F_{\text{opt}} + \text{Opt} \leq 4F_{\text{opt}}$ (as we can take the original schedule and delay every job by $3F_{\text{opt}}$). To prove theorem 4, we will in fact prove the following stronger result: Double-fit never fails for any instance where the restricted Opt has value at most $4F_{\text{opt}}$.

The Invariants

Fix an epoch k. Let $A_i(k)$ and $B_i(k)$ denote the total work on machines $[i : m]$ in Double-fit's schedule just before and just after all the jobs from interval I_k are scheduled respectively. Similarly, let $A_i^{\text{opt}}(k)$ and $B_i^{\text{opt}}(k)$ be the total work remaining on machines $[i : m]$ in Opt's schedule.

We will show that the following two invariants hold at each epoch k.

$$A_i(k) \leq A_i^{\text{opt}}(k) + F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}. \tag{1}$$

$$B_i(k) \leq \max\left\{3F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}, B_i^{\text{opt}}(k)\right\} + F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}. \tag{2}$$

Roughly speaking, invariants (1) and (2) show that the load on any suffix of Double-fit's machines stays close to Opt's load on those machines, both before and after the jobs are scheduled in epoch k. We will prove that (1) and (2) hold by a careful induction over i and k.

Before we prove these invariants, let us first see why they imply Theorem 4.

Proof of Theorem 4. Consider a fixed epoch k. As the (restricted) Opt has maximum flow-time at most $4F_{\text{opt}}$, for each i it must hold that $B_i^{\text{opt}}(k) \leq 4F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}$. Thus by (2) it follows that $B_i^{\text{opt}}(k) \leq 5F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}$ for each i. Choosing $i = m$, this implies that

Double-fit never loads machine m above $5F_{\text{opt}}$ and thus never fails (as machine m always has room for an additional job). ◀

Proving the Invariants

The strategy for proving that (1) and (2) hold at all epochs k will be to show the following two lemmas.

▶ **Lemma 5.** *If at epoch k,* (1) *holds for all machines, then* (2) *also holds for all machines.*

The next step will be to relate the conditions at epochs k and $k+1$.

▶ **Lemma 6.** *If at any epoch k,* (2) *holds for all machines, then* (1) *also holds for all machines at epoch $k+1$.*

As (1) trivially holds for $k = 0$ (as $A_i(0) = A_i^{\text{opt}}(0) = 0$ for all i), applying Lemma 5 and Lemma 6 alternately implies that (1) and (2) hold for all k.

3.2 Proof of Lemma 5

We first show that Double-fit is conservative in scheduling small jobs on fast machines.

▶ **Lemma 7.** *Let $i_1 < i_2$. If some job j of class i_1 is scheduled by* Double-fit *onto machine i_2 during the saturation phase (i.e. using threshold $3F_{opt}$), then* all *jobs of class i for $i_1 < i \leq i_2$ are also scheduled during the saturation phase.*

Proof. Consider the state of Double-fit's machines just before j was scheduled. As j is scheduled on machine i_2 during the saturation phase, the load on i_2 must be below $3F_{\text{opt}}$ at that point. As jobs of class i for $i_1 < i \leq i_2$ were considered before class i_1-jobs, the load on i_2 was also below $3F_{\text{opt}}$ after scheduling class i jobs, and thus Double-fit must have never switched to the Slow-fit phase while considering class i. ◀

Next we define the notion of *separated* machines, which will play a crucial role in the analysis.

▶ **Definition 8.** Machines i_1 and i_2 ($i_1 < i_2$) are *separated* at epoch k if Double-fit scheduled no jobs from classes $[1 : i_1]$ onto machines $[i_2 : m]$ at epoch k.

The following lemma shows that if two consecutive machines are separated, it is easy to relate epochs k and $k+1$.

▶ **Lemma 9.** *If machines $i-1$ and i are separated at epoch k, then* (1) *implies* (2) *for machine i. Moreover this trivially holds for machine $i = 1$.*

Proof. As machines $i-1$ and i are separated at epoch k, no jobs from class $[1 : i-1]$ were scheduled onto machines $[i : m]$ at epoch k. Thus

$$B_i(k) = A_i(k) + \sum_{i'=i}^{m} |J_{i'}|, \tag{3}$$

where $|J_i|$ represents the total work of all jobs in J_i.

As jobs from J_i cannot be scheduled onto machines $[1 : i-1]$ in an optimal schedule, we also obtain

$$B_i^{\text{opt}}(k) \geq A_i^{\text{opt}}(k) + \sum_{i'=i}^{m} |J_i|. \tag{4}$$

This implies that

$$B_i(k) = A_i(k) + \sum_{i'=i}^{m} |J_{i'}| \le A_i(k) + B_i^{\text{opt}}(k) - A_i^{\text{opt}}(k) \le B_i^{\text{opt}}(k) + F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}, \tag{5}$$

where the last step follows by our assumption that (1) holds for (i,k).

Finally for $i = 1$, we observe that both (3) and (4) hold with equality, and hence the result holds trivially. ◀

We now have all the tools we need to prove Lemma 5.

Proof of Lemma 5. We use induction over i in the order of larger to smaller i. In particular, to prove that (2) holds for some pair (i,k), we assume that (1) holds for all (i',k) and that (2) holds for all (i',k) with $i' > i$. As the base case note that this is vacuously true for $i = m+1$ (as all relevant quantities are 0).

We consider three cases depending on how Double-fit assigns jobs from classes $[1 : i-1]$ to machines $[i : m]$.

1. *No jobs from class* $[1 : i-1]$ *were scheduled onto machines* $[i : m]$*:* In this case, machines $i-1$ and i are separated and (2) follows from Lemma 9.

2. *Jobs from classes* $[1 : i-1]$ *are only scheduled onto machines* $[i : m]$ *during the saturation phase:* Let $i_{\max} \ge i$ denote the smallest index such that machines $i-1$ and $i_{\max}+1$ are separated (if no such machine exists, set $i_{\max} = m$). By the inductive hypothesis, we can assume that (2) holds for $i_{\max}+1$. In the case where $i_{\max} = m$, this holds vacuously. As jobs from classes $[1 : i-1]$ are assigned to $[i : m]$ (and hence to $i_{\max}$) during the saturation phase, Lemma 7 implies that all jobs in classes $[i : i_{\max}]$ were also scheduled during the saturation phase, which implies that all machines $[i : i_{\max}]$ are loaded below $4F_{\text{opt}}$. This gives us the following:

$$\begin{aligned} B_i(k) &\le 4F_{\text{opt}} \sum_{i'=i}^{i_{\max}} s_{i'} + B_{i_{\max}+1}(k) \\ &\le 4F_{\text{opt}} \sum_{i'=i}^{i_{\max}} s_{i'} + \max\left\{3F_{\text{opt}} \sum_{i'=i_{\max}+1}^{m} s_{i'}, B_{i_{\max}+1}^{\text{opt}}(k)\right\} + F_{\text{opt}} \sum_{i'=i_{max}+1}^{m} s_{i'} \\ &= 3F_{\text{opt}} \sum_{i'=i}^{i_{\max}} s_{i'} + \max\left\{3F_{\text{opt}} \sum_{i'=i_{\max}+1}^{m} s_{i'}, B_{i_{\max}+1}^{\text{opt}}(k)\right\} + F_{\text{opt}} \sum_{i'=i}^{m} s_{i'} \\ &\le \max\left\{3F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}, B_i^{\text{opt}}(k)\right\} + F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}, \end{aligned}$$

   where the second inequality follows from the inductive hypothesis for machine $i_{\max}+1$.

3. *Some job* j *from class* $[1 : i-1]$ *was scheduled onto machines* $[i : m]$ *during Slow-fit phase (using threshold* $6F_{opt}$*):* We assume that $i > 1$, otherwise the result follows from case 1. Let $i_{min} < i$ denote the largest index such that machines $[i_{min} : i-1]$ have load more than $5F_{\text{opt}}$ and machine $i_{\min}-1$ has load at most $5F_{\text{opt}}$. If no such machine exists, set $i_{\min} = 1$. $i_{\min}$ is well-defined as $i > 1$ and machine $i-1$ must have load more than $5F_{\text{opt}}$ as job j from class $[1 : i-1]$ was assigned to a machine in $[i : m]$ during the Slow-fit phase.

▶ **Claim 1.** *Machines* $i_{\min}-1$ *and* $i_{\min}$ *are separated or* $i_{\min} = 1$.

Proof. This is trivially true if $i_{\min} = 1$.
If $i_{\min} > 1$, suppose that some job j' from class $[1 : i_{\min} - 1]$ was scheduled onto machines $[i_{\min} : m]$. Now j' cannot be scheduled during the Slow-fit phase as this would imply that the load on $i_{\min} - 1$ was more than $5F_{\text{opt}}$, which contradicts the choice of $i_{\min}$.
So all jobs in $[1 : i_{\min} - 1]$ that were assigned to $[i_{\min} : m]$ must have been assigned during the saturation phase. Let $i' \geq i_{\min}$ denote the largest index where such a job is assigned. By Lemma 7, it must be that all machines $[i_{\min} : i']$ were assigned load during the saturation phase and must have load at most $4F_{\text{opt}}$. This contradicts that $i_{\min}$ has load more than $5F_{\text{opt}}$. ◀

By Lemma 9 applied to $i_{\min}$, we get that (2) holds for machine $i_{\min}$ and thus

$$B_{i_{\min}}(k) \leq \max\left\{3F_{\text{opt}} \sum_{i'=i_{\min}}^{m} s_{i'}, B^{\text{opt}}_{i_{\min}}(k)\right\} + F_{\text{opt}} \sum_{i'=i_{\min}}^{m} s_{i'}. \tag{6}$$

Furthermore, by choice of $i_{\min}$ all the machines in $[i_{\min} : i - 1]$ are loaded above $5F_{\text{opt}}$. This implies that

$$B_i(k) \leq B_{i_{\min}}(k) - 5F_{\text{opt}} \sum_{i'=i_{\min}}^{i-1} s_{i'}. \tag{7}$$

As every machine is loaded below $4F_{\text{opt}}$ in an optimal schedule, we also have

$$B^{\text{opt}}_{i_{\min}}(k) \leq B^{\text{opt}}_{i}(k) + 4F_{\text{opt}} \sum_{i'=i_{\min}}^{i-1} s_{i'}. \tag{8}$$

Adding (6) and (7) we obtain that

$$\begin{aligned}
B_i(k) &\leq \max\left\{3F_{\text{opt}} \sum_{i'=i_{\min}}^{m} s_{i'}, B^{\text{opt}}_{i_{\min}}(k)\right\} + F_{\text{opt}} \sum_{i'=i_{\min}}^{m} s_{i'} - 5F_{\text{opt}} \sum_{i'=i_{\min}}^{i-1} s_{i'} \\
&\leq \max\left\{4F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}, B^{\text{opt}}_{i_{\min}}(k) + F_{\text{opt}} \sum_{i'=i}^{m} s_{i'} - 4F_{\text{opt}} \sum_{i'=i_{\min}}^{i-1} s_{i'}\right\} \\
&\leq \max\left\{4F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}, B^{\text{opt}}_{i}(k) + F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}\right\} \qquad \text{By (8)} \\
&= \max\left\{3F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}, B^{\text{opt}}_{i}(k)\right\} + F_{\text{opt}} \sum_{i'=i}^{m} s_{i'},
\end{aligned}$$

which implies that (2) holds for i.

◀

3.3 Proof of Lemma 6

We now prove Lemma 6, which is relatively easier.

Proof of Lemma 6. We will apply induction over i (in decreasing order of machines). Consider epoch k. We assume that (2) holds for all i' at epoch k, and that (1) holds for all $i' > i$ at epoch $k + 1$. For the base case of $i = m + 1$ the lemma follows trivially since all the relevant quantities are 0.

Consider some machine i. We consider two cases depending on the load of machine i after the jobs were scheduled at epoch k.

1. *Machine i has load at most $4F_{opt}$ after epoch k, i.e., $B_i(k) - B_{i+1}(k) \leq 4F_{opt} \cdot s_i$:* At epoch $k+1$ before the jobs arriving during interval I_{k+1} are scheduled, the load of machine i will be at most F_{opt}. Thus we have that
$$\begin{aligned} A_i(k+1) &\leq A_{i+1}(k+1) + F_{\text{opt}} \cdot s_i \\ &\leq A^{\text{opt}}_{i+1}(k+1) + F_{\text{opt}} \sum_{i'=i+1}^{m} s_{i'} + F_{\text{opt}} \cdot s_i \\ &\leq A^{\text{opt}}_{i}(k+1) + F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}. \end{aligned}$$
Here the second inequality follows by the inductive hypothesis for machine $i+1$, and the third inequality follows as $A_i^{\text{opt}}(k+1)$ is non-decreasing as i decreases.
2. *Machine i is loaded above $4F_{opt}$ after epoch k, i.e., $B_i(k) - B_{i+1}(k) > 4F_{opt} \cdot s_i$:* In this case, some job j must have been scheduled onto machine i during the Slow-fit phase. This only happens if j could not be scheduled during the saturation phase. In particular, this implies that all the machines $[i : m]$ (which is surely a subset of machines where j could have been scheduled) were loaded above $3F_{\text{opt}}$. So the total work on all machines $[i : m]$ decreases by exactly $3F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}$ during interval I_{k+1}.
Thus we have that
$$A_i(k+1) = B_i(k) - 3F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}. \tag{9}$$
Similarly, as Opt can complete at most $3F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}$ on machines $[i : m]$ during this interval, we have
$$A_i^{\text{opt}}(k+1) \geq B_i^{\text{opt}}(k) - 3F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}. \tag{10}$$
As (2) holds for each i at epoch k, we obtain that
$$\begin{aligned} A_i(k+1) &\leq B_i(k) - 3F_{\text{opt}} \sum_{i'=i}^{m} s_{i'} \\ &\leq B_i^{\text{opt}}(k) + F_{\text{opt}} \sum_{i'=i}^{m} s_{i'} - 3F_{\text{opt}} \sum_{i'=i}^{m} s_{i'} \qquad \text{By (2)} \\ &\leq A^{\text{opt}}_{i}(k+1) + F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}, \end{aligned}$$
and hence (1) holds for i at epoch $k+1$, which completes the proof. ◀

3.4 Removing the assumption of knowledge of Opt

We describe a variant of the standard doubling trick where we increase the online estimate of Opt by only 1.5 times at each step.

Consider some epoch k where the algorithm first fails with the current guess of F_{opt}. It must be that (2) does not hold. In particular, (1) holds at epoch k as (2) holds at $k-1$. Now, Lemma 5 implies that $F_{\text{opt}} < \text{Opt}$. We then abort epoch k, and do not schedule any jobs. Instead, we set $F'_{\text{opt}} = 1.5F_{\text{opt}}$ and redefine the new epoch to be the time $(k-1)F_{\text{opt}} + 3F'_{\text{opt}}$. Note that between these epochs $4.5F_{\text{opt}}$ time passes, so at the next epoch the load on all machines in the schedule of DOUBLE-FIT will be at most $6F_{\text{opt}} - 3F'_{\text{opt}} = 1.5F_{\text{opt}} = F'_{\text{opt}}$.

This implies that for all i

$$A_i(k) \leq F_{\text{opt}} \sum_{i'=i}^{m} s_{i'} \leq A_i^{\text{opt}}(k) + F_{\text{opt}} \sum_{i'=i}^{m} s_{i'}.$$

The crucial point is that (1) holds for all machines i at this new epoch irrespective of the workload of the new restricted Opt (with parameter F'_{opt}). Thus, (2) holds if $F_{\text{opt}} \geq \text{Opt}$ and Double-fit proceeds as normal.

4 Other Lower bounds

We also show simple (but strong) lower bounds for weighted maximum flow-time and maximum stretch.

▶ **Lemma 10.** *Any algorithm for minimizing maximum weighted flow-time on identical machines must have a competitive ratio of $\Omega(W)$, where W is the ratio between the largest and smallest weight.*

Proof. Consider the following instance on 2 machines. At time $t = 0$ we receive 2 jobs of size w with weight 1. Now, any algorithm has three options: (i) it schedules both jobs on the same machine, (ii) it schedules both jobs on different machines, or (iii) it waits before assigning jobs. In all three cases, we show that it will end up trailing by at least w work behind an optimal schedule.

In option (i), w work remains at time $t = w$, while optimum is empty. In option (ii) we instantly receive another $2w$-sized job with weight 1, so that one of our machines has load $3w$. At time $t = 2w$ we have w load remaining while an optimal schedule is empty. In option (iii) we receive no jobs until algorithm decide to choose (i) or (ii). If we have not chosen by time $t = w$ we are trailing $2w$ work behind an optimal schedule.

Once we trail w work behind optimum, at every unit time step we receive 2 unit-size jobs of weight w. If the trailing jobs are ever to be finished, at least $w/2$ delay is incurred on the weight w jobs, implying an objective value of $\Omega(w^2)$. Opt on other hand has value $O(w)$. ◀

A lower bound of $\Omega(W^{0.4})$ for maximum weighted flow-time follows from [10], using the analogy between delay factor and weighted maximum flow-time described in [3]. By replacing the unit size jobs by unit weight jobs in the above lower bound instance, this also directly implies an $\Omega(S)$ lower bound on the competitive ratio for maximum stretch [3] where S is the ratio between the size of the largest and the smallest job.

5 Concluding Remarks

Note that our algorithm Double-fit is not immediate dispatch, i.e., it does not dispatch a job to a machine immediately upon arrival. We are unable to extend the ideas here to obtain an $O(1)$-competitive immediate dispatch algorithm, and it is not clear to us whether such an algorithm exists. Given that in the unrelated setting, there can be no $O(1)$-speed, $O(1)$-competitive immediate dispatch algorithm [3] (while there is a $(1+\epsilon)$-speed, $O(1/\epsilon)$-competitive algorithm), it would be quite interesting to resolve this question.

References

1 Susanne Albers. *Introduction to scheduling*, chapter Online scheduling, pages 51–73. Chapman and Hall/CRC, 2010.

2 Christoph Ambühl and Monaldo Mastrolilli. On-line scheduling to minimize max flow time: an optimal preemptive algorithm. *Oper. Res. Lett.*, 33(6):597–602, 2005.

3 S. Anand, Karl Bringmann, Tobias Friedrich, Naveen Garg, and Amit Kumar. Minimizing maximum (weighted) flow-time on related and unrelated machines. In *ICALP (1)*, pages 13–24, 2013.

4 S. Anand, Naveen Garg, and Amit Kumar. Resource augmentation for weighted flow-time explained by dual fitting. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1228–1241, 2012.

5 Yossi Azar. On-line load balancing. In Amos Fiat and Gerhard J. Woeginger, editors, *Online Algorithms*, volume 1442 of *Lecture Notes in Computer Science*, pages 178–195. Springer, 1998.

6 Yossi Azar, Bala Kalyanasundaram, Serge A. Plotkin, Kirk Pruhs, and Orli Waarts. On-line load balancing of temporary tasks. *J. Algorithms*, 22(1):93–110, 1997.

7 Nikhil Bansal and Janardhan Kulkarni. Minimizing flow-time on unrelated machines. In *Symposium on Theory of Computing, STOC*, 2015, to appear.

8 Michael A. Bender, Soumen Chakrabarti, and S. Muthukrishnan. Flow and stretch metrics for scheduling continuous job streams. In *Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 270–279, 1998.

9 Niv Buchbinder and Joseph Naor. The design of competitive online algorithms via a primal-dual approach. *Foundations and Trends in Theoretical Computer Science*, 3(2-3):93–263, 2009.

10 Chandra Chekuri and Benjamin Moseley. Online scheduling to minimize the maximum delay factor. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1116–1125. Society for Industrial and Applied Mathematics, 2009.

11 Yookun Cho and Sartaj Sahni. Bounds for list schedules on uniform processors. *SIAM Journal on Computing*, 9(1):91–103, 1980.

12 Naveen Garg. Minimizing average flow-time. In *Efficient Algorithms, Essays Dedicated to Kurt Mehlhorn on the Occasion of His 60th Birthday*, pages 187–198, 2009.

13 Sungjin Im, Benjamin Moseley, and Kirk Pruhs. A tutorial on amortized local competitiveness in online scheduling. *SIGACT News*, 42(2):83–97, 2011.

14 Bala Kalyanasundaram and Kirk Pruhs. Speed is as powerful as clairvoyance. *J. ACM*, 47(4):617–643, 2000.

15 Kirk Pruhs, Jiri Sgall, and Eric Torng. *Handbook of Scheduling: Algorithms, Models, and Performance Analysis*, chapter Online Scheduling. CRC Press, 2004.