The Lazy Bureaucrat Scheduling Problem

Esther M. Arkin[1]*, Michael A. Bender[2]**, Joseph S. B. Mitchell[1]***, and Steven S. Skiena[2]†

[1] Department of Applied Mathematics and Statistics
[2] Department of Computer Science
State University of New York, Stony Brook, NY 11794, USA

Abstract. We introduce a new class of scheduling problems in which the optimization is performed by the worker (single "machine") who performs the tasks. The worker's objective may be to minimize the amount of work he does (he is "lazy"). He is subject to a constraint that he must be busy when there is work that he *can* do; we make this notion precise, particularly when preemption is allowed. The resulting class of "perverse" scheduling problems, which we term "Lazy Bureaucrat Problems," gives rise to a rich set of new questions that explore the distinction between maximization and minimization in computing optimal schedules.

1 Introduction

Scheduling problems have been studied extensively from the point of view of the objectives of the enterprise that stands to gain from the completion of the set of jobs. We take a new look at the problem from the point of view of the workers who perform the tasks that earn the company its profits. In fact, it is natural to expect that some employees may lack the motivation to perform at their peak levels of efficiency, either because they have no stake in the company's profits or because they are simply lazy. The following example illustrates the situation facing a "typical" office worker, who may be one small cog in a large bureaucracy:

Example. It is 3:00 p.m., and Dilbert goes home at 5:00 p.m. Dilbert has two tasks that have been given to him: one requires 10 minutes, the other requires an hour. If there is a task in his "in-box," Dilbert must work on it, or risk getting fired. However, if he has multiple tasks, Dilbert has the freedom to choose which one to do first. He also knows that at 3:15, another task will appear — a 45-minute personnel meeting. If Dilbert begins the 10-minute task first, he will be free to attend the personnel meeting at 3:15 and then work on the hour-long task

* estie@ams.sunysb.edu; Partially supported by NSF (CCR-9732220).
** bender@cs.sunysb.edu.
*** jsbm@ams.sunysb.edu; Partially supported by Boeing, Bridgeport Machines, Sandia National Labs, Seagull Technologies, Sun Microsystems, and NSF (CCR-9732220).
† skiena@cs.sunysb.edu; Partially supported by NSF (CCR-9625669), and ONR (431-0857A).

from 4:00 until 5:00. On the other hand, if Dilbert is part way into the hour-long job at 3:15, he may be excused from the meeting. After finishing the 10-minute job by 4:10, he will have 50 minutes to twiddle his thumbs, iron his tie, or enjoy engaging in other mindless trivia. Naturally, Dilbert prefers this latter option.

There is also an historical example of an actual situation in which it proved crucial to schedule tasks inefficiently, as documented in the book/movie *Schindler's List* [4]. It was essential for the workers and management of Schindler's factory to appear to be busy at all times, in order to stay in operation, but they simultaneously sought to minimize their contribution to the German war effort.

These examples illustrate a general and natural type of scheduling problem, which we term the "Lazy Bureaucrat Problem" (LBP), in which the goal is to schedule jobs as *inefficiently* (in some sense) as possible. We propose that these problems provide an interesting set of algorithmic questions, which may also lead to discovery of structure in traditional scheduling problems. (Several other combinatorial optimization problems have been studied "in reverse," leading, e.g., to maximum TSP, maximum cut, and longest path; such inquiries can lead to better understanding of the structure and algorithmic complexity of the original optimization problem.) Our investigations may also be motivated by a "game theoretic" view of the employee-employer system.

1.1 The Model

There is a vast literature on a variety of scheduling problems; see, e.g., some of the recent surveys [3, 5, 7]. Here, we consider a set of jobs $1 \ldots n$ having processing times (lengths) $t_1 \ldots t_n$ respectively. Job i arrives at time a_i and has its deadline at time d_i. We assume throughout this paper that t_i, a_i, and d_i have nonnegative integral values. The jobs have *hard deadlines,* meaning that each job i can only be executed during its allowed interval $I_i = [a_i, d_i]$; we also call I_i the job's *window*. We let $c_i = d_i - t_i$ denote the *critical time* of job i; job i must be started by time c_i if there is going to be any chance of completing it on time.

The jobs are executed on a single processor, the (lazy) *bureaucrat*. The bureaucrat executes only one job at a time. (We leave the case of multiple processors for future work.)

Greedy Requirement. The bureaucrat chooses a subset of jobs to execute. Since his goal is to minimize his effort, he prefers to remain idle all the time and to leave all the jobs unexecuted. However, this scenario is forbidden by what we call the *greedy requirement,* which requires that the bureaucrat work on an executable job, if there are any executable jobs. A job is "executable" if it has arrived, its deadline has not yet passed, and it is not yet fully processed. In the case with preemption, there may be other constraints that govern whether or not a job is executable; see Section 3.

Objective Functions. In traditional scheduling problems, if it is impossible to complete the set of all jobs by their deadlines, one typically tries to optimize according to some objective, e.g., to maximize a weighted sum of on-time jobs,

to minimize the maximum lateness of the jobs, or to minimize the number of late jobs. For the LBP we consider three different objective functions, which naturally arise from considering the bureaucrat's goal of being inefficient:

(1) *Minimize the total amount of time spent working.* This objective naturally appeals to a "lazy" bureaucrat.

(2) *Minimize the weighted sum of completed jobs.* Here, we usually assume that the weight of job i is its length, t_i; however, other weights (e.g., unit weights) are also of interest. This objective appeals to a "spiteful" bureaucrat whose goal it is to minimize the fees that the company collects on the basis of his labors, assuming that the fee (in proportion to the task length, or a fixed fee per task) is collected only for those tasks that are actually completed.

(3) *Minimize the* makespan, *the maximum completion time of the jobs.* This objective appeals to an "impatient" bureaucrat, whose goal it is to go home as early as possible, at the completion of the last job he is able to complete. He cares about the number of hours spent at the office, not the number of hours spent doing work (productive or otherwise) at the office.
Note that, in contrast with standard scheduling problems, the makespan in the LBP changes; it is a function of which jobs have passed their deadlines and can no longer be executed.

Additional Parameters of the Model. As with most scheduling problems, additional parameters of the model must be set. For example, one must explicitly allow or forbid *preemption* of jobs. If a job is preempted, it is interrupted and may be resumed later at no additional cost. If preemption is forbidden, then once a job is begun, it must be completed without interruptions.

One must also specify whether scheduling occurs *on-line* or *off-line.* A scheduling algorithm is considered to be off-line if all the jobs are known to the scheduler at the outset; it is on-line if the jobs are known to the scheduler only as they arrive. In this paper we restrict ourselves to off-line scheduling.

1.2 Our Results

In this paper, we introduce the Lazy Bureaucrat Problem and develop algorithms and hardness results for several versions of it. From these results, we derive some general characteristics of this new class of scheduling problems and describe (1) situations in which traditional scheduling algorithms extend to our problems and (2) situations in which these algorithms no longer apply.

No Preemption. We prove that the LBP is NP-complete, as is generally the case for traditional scheduling problems. Thus, we focus on special cases to study algorithms. When all jobs have unit size, optimal schedules can be found in polynomial time. The following three cases have pseudo-polynomial algorithms: (1) when each job i's interval I_i is less than twice the length of i; (2) when the ratios of interval length to job length and longest job to shortest job are both bounded; and (3) when all jobs arrive in the system at the same time. These

last scheduling problems are solved using dynamic programming both for Lazy Bureaucrat and traditional metrics. Thus, in these settings, the Lazy Bureaucrat metrics and traditional metrics are solved using similar techniques.

From the point of view of approximation, however, the standard and Lazy Bureaucrat metrics behave differently. Standard metrics typically allow polynomial-time algorithms having good approximation ratios, whereas we show that the Lazy Bureaucrat metrics are difficult to approximate. This hardness derives more from the greedy requirement and less from the particular metric in question. The greedy requirement appears to render the problem substantially more difficult. (Ironically, even in standard optimization problems, the management often tries to impose this requirement, because it naively appears to be desirable.)

Preemption. The greedy requirement dictates that the worker must stay busy while work is in the system. If the model allows preemption we must specify under what conditions a job can be interrupted or resumed. We distinguish three versions of the preemption rules, which we list from most permissive to most restrictive. In particular possible constraints on what the worker can execute include (I) any job that has arrived and is before its deadline, (II) any job that has arrived and for which there is still time to complete it before its deadline, or (III) any job that has arrived, but with the constraint that if it is started, it must eventually be completed.

We consider all three metrics and all three versions of preemption. We show that, for all three metrics, version I is polynomially solvable, and version III is NP-complete. Many of the hardness results for no preemption carry over to version III. However, the question of whether the problem is strongly NP-complete remains open.

Our main results are for version II. We show that the general problem is NP-complete. Then, we focus on minimizing the makespan in two complementary special cases: (1) All jobs have a common arrival time and arbitrary deadlines; (2) All jobs have a common deadline and arbitrary arrival times. We show that the first problem is NP-complete, whereas the second problem can be solved in polynomial time. These last results illustrate a curious feature of the LBP. One can convert one special case into the other by reversing the direction of time. In the LBP, unlike many scheduling settings, this reversing of time changes the complexity of the problem.

Due to space limitations here, many proofs are deferred to the full paper[1].

2 LBP: No Preemption

In this section, we assume that no job can be preempted: if a job is started, then it is performed without interruption until it completes. We show that the Lazy Bureaucrat Problem (LBP) without preemption is strongly NP-complete and is not approximable to within any factor. These hardness results distinguish

[1] See `http://www.ams.sunysb.edu/~jsbm/lazy/paper.ps.gz`.

our problem from traditional scheduling metrics, which can be approximated in polynomial time, as shown in the recent paper of [1]. We show, however, that several special cases of the problem have pseudo-polynomial time algorithms, using applications of dynamic programming.

2.1 Hardness Results

We begin by describing the relationship between the three different objective functions in the case of no preemption. The problem of minimizing the total work is a special case of the problem of minimizing the weighted sum of completed jobs, because every job that is executed must be completed. (The weights become the job lengths.) Furthermore, if all jobs have the same arrival time, say time zero, then the two objectives, minimizing total work and minimizing makespan (go home early) are equivalent, since no feasible schedule will have any gaps. Our first hardness theorem applies therefore to all three objective functions:

Theorem 1. *The Lazy Bureaucrat Problem with no preemption is (weakly) NP-complete, and is not approximable to within any fixed factor, even when arrival times are all the same.*

Proof. We use a reduction from the SUBSET SUM problem [2]. □

As we show in Section 2.2, the problem from Theorem 1 has a pseudopolynomial-time algorithm. However, if arrival times and deadlines are arbitrary integers, the problem is strongly NP-complete. The given reduction applies to all three objective functions.

Theorem 2. *The Lazy Bureaucrat Problem with no preemption is strongly NP-complete, and is not approximable to within any fixed factor.*

Proof. We use a reduction from the 3-PARTITION problem [2]. □

2.2 Algorithms for Special Cases

Unit-Length Jobs. Consider the special case of the LBP in which all jobs have unit processing times. (Recall that all inputs are assumed to be integral.) The Latest Deadline First (LDF) scheduling policy selects the job in the system having the latest deadline.

Theorem 3. *The Latest Deadline First (LDF) scheduling policy minimizes the amount of executed work.*

Narrow Windows. Consider now the version in which jobs are large in comparison with their intervals, that is, the intervals are "narrow." Let R be a bound on the ratio of window length to job length; i.e., for each job i, $d_i - a_i < R \cdot t_i$. We show that a pseudo-polynomial algorithm exists for the case of sufficiently narrow windows, that is, when $R \leq 2$.

Lemma 1. *Assume that for each job i, $d_i - a_i < 2t_i$. Then, if job i can be scheduled before job j, then job j cannot be scheduled before job i.*

Under this assumption, the ordering of any subset of jobs in a schedule is uniquely determined, allowing us to use dynamic programming to solve the problem:

Theorem 4. *Suppose that for each job i, $d_i - a_i < 2t_i$. Let $K = \max_i d_i$. Then the LBP can be solved in $O(nK \max(n, K))$ time.*

Note that for $R > 2$ we know of no efficient algorithm without additional conditions. Let W be a bound on the ratio of longest window to shortest window, and let Δ be a bound on the ratio of the longest job to the shortest job. Note that bounds on R and Δ imply a bound on W, and bounds on R and W imply a bound on Δ. However, a bound on Δ alone is not sufficient for a pseudopolynomial time algorithm.

Theorem 5. *Even with a bound on the ratio Δ, the LBP with no preemption is strongly NP-complete. It cannot be approximated to within a factor of $\Delta - \epsilon$, for any $\epsilon > 0$, unless P=NP.*

Proof. Modify the reduction from 3-partition of Theorem 2. □

Bounds on both Δ and R are sufficient to yield a pseudo-polynomial algorithm:

Theorem 6. *Let $K = \max_i d_i$. Given bounds on R and Δ, the Lazy Bureaucrat Problem with no preemption can be solved in $O(K \cdot n^{4R \lg \Delta})$.*

Proof. We modify the dynamic programming algorithm of Theorem 4. □

Jobs Having a Common Release Time. In the next version of the problem all jobs are released at time zero, i.e., $a_i = 0$ for all i. This problem can be solved in polynomial time by dynamic programming. The dynamic programming works because of the following structural result: There exists an optimal schedule that executes the jobs Earliest Due Date (EDD).

In fact this problem is a special case of the following general problem: Minimizing the weighted sum of jobs not completed by their deadlines. This problem was solved by [6], using the same structural result.

Theorem 7. *The LBP can be solved in pseudo-polynomial time when all jobs have a common release time.*

3 LBP: Preemption

In this section we consider the Lazy Bureaucrat Problem in which jobs may be preempted: a job in progress can be set aside, while another job is processed, and then possibly resumed later. It is important to distinguish among different constraints that specify which jobs are available to be processed. We consider three natural choices of such constraints:

Constraint I: In order to work on job i at time τ, we require only that the current time τ lies within the job's interval I_i: $a_i \leq \tau \leq d_i$.

Constraint II: In order to work on job i at time τ, we require not only that the current time τ lies within the job's interval I_i, but also that the job has a *chance* to be completed, e.g., if it is processed without interruption until completion.
This condition is equivalent to requiring that $\tau \leq c'_i$, where $c'_i = d_i - t_i + y_i$ is the *adjusted critical time* of job i: c'_i is the latest possible time to start job i, in order to meet its deadline d_i, given that an amount y_i of the job has already been completed.

Constraint III: In order to work on job i, we require that $\tau \in I_i$. Further, we require that any job that is started is eventually completed.

As before, we consider the three objective functions (1)–(3), in which the goal is to minimize the total time working (regardless of which jobs are completed), the weighted sum of completed jobs, or the makespan of the schedule (the "go home" time).

The third constraint makes the problem with preemption quite similar to the one with no preemption. In fact, if all jobs arrive at the same time ($a_i = 0$ for all i), then the three objective functions are equivalent, and the problem is hard:

Theorem 8. *The LBP with preemption, under constraint III (one must complete any job that is begun), is NP-complete and hard to approximate.*

3.1 Minimizing Total Time Working

Theorem 9. *The LBP with preemption, under constraint I (one can work on any job in its interval) and objective (1) (minimize total time working), is polynomially solvable.*

Proof. The algorithm schedules jobs according to latest deadline first (LDF). □

Theorem 10. *The LBP with preemption, under constraint II (one can only work on jobs that can be completed) and objective (1) (minimize total time working), is (weakly) NP-complete.*

Proof. If all arrival times are the same, then this problem is equivalent to the one in which the objective function is minimize the makespan, which is shown to be NP-complete in Theorem 14. □

3.2 Minimizing Weighted Sum of Completed Jobs

Theorem 11. *The LBP with preemption, under constraint I (one can work on any job in its interval) and objective (2) (minimize the weighted sum of completed jobs), is polynomially solvable.*

Theorem 12. *The LBP with preemption, under constraint II (one can only work on jobs that can be completed) and objective (2) (minimize the weighted sum of completed jobs), is (weakly) NP-complete.*

3.3 Minimizing Makespan: Going Home Early

We assume now that the bureaucrat's goal is to go home as soon as possible.

We begin by noting that if the arrival times are all the same ($a_i = 0$, for all i), then the objective (3) (go home as soon as possible) is in fact equivalent to the objective (1) (minimize total time working), since, under any of the three constraints I–III, the bureaucrat will be busy nonstop until he can go home.

We note, however, that if the *deadlines* are all the same ($d_i = D$, for all i), then the objectives (1) and (3) are quite different. Consider the following example. Job 1 arrives at time $a_1 = 0$ and is of length $t_1 = 2$, job 2 arrives at time $a_2 = 0$ and is of length $t_2 = 9$, job 3 arrives at time $a_3 = 8$ and is of length $t_3 = 2$, and all jobs have deadline $d_1 = d_2 = d_3 = 10$. Then, in order to minimize total time working, the bureaucrat will do jobs 1 and 3, a total of 4 units of work, and will go home at time 10. However, in order to go home as soon as possible, the bureaucrat will do job 2, performing 9 units of work, and go home at time 9 (since there is not enough time to do either job 1 or job 3).

Theorem 13. *The LBP with preemption, under constraint I (one can do any job in its interval) and objective (3) (go home as early as possible), is polynomially solvable.*

Proof. The algorithm is to schedule by latest deadline first (LDF). The proof is similar to the one given in Theorem 9. □

If instead of constraint I we impose constraint II, the problem becomes hard:

Theorem 14. *The LBP with preemption, under constraint II (one can only work on jobs that can be completed) and objective (3) (go home as early as possible), is (weakly) NP-complete, even if all arrival times are the same.*

Proof. We give a reduction from SUBSET SUM. Consider an instance of SUBSET SUM given by a set S of n positive integers, $x_1, x_2, \ldots, x_n$, and target sum T. We construct an instance of the required version of the LBP as follows. For each integer x_i, we have a job i that arrives at time $a_i = 0$, has length $t_i = x_i$, and is due at time $d_i = T + x_i - \epsilon$, where ϵ is a small constant (it suffices to use $\epsilon = \frac{n}{3}$). In addition, we have a "long" job $n+1$, with length $t_{n+1} > T$, that arrives at time $a_{n+1} = 0$ and is due at time $d_{n+1} = T - 2\epsilon + t_{n+1}$. We claim that it is possible for the bureaucrat to go home by time T if and only if there exists a subset of $\{x_1, \ldots, x_n\}$ that sums to exactly T.

If there is a subset of $\{x_1, \ldots, x_n\}$ that sums to exactly T, then the bureaucrat can perform the corresponding subset of jobs (of total length T) and go home at time T; he is able to avoid doing any of the other jobs, since their critical times fall at an earlier time ($T - \epsilon$ or $T - 2\epsilon$), making it infeasible to begin them at time T, by our assumption.

If, on the other hand, the bureaucrat is able to go home at time T, then we know the following (details omitted in this abstract):

(a) He must have just completed a job at time T.

(b) He must have been busy the entire time from 0 until time T.
(c) If he starts a job, then he must finish it.

We conclude that the bureaucrat must complete a set of jobs whose lengths sum exactly to T. Thus, we have reduced SUBSET SUM to our problem, showing that it is (weakly) NP-complete. □

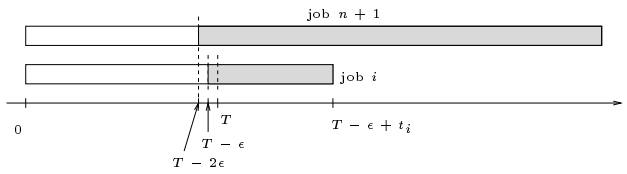

Fig. 1. Proof of hardness of LBP with preemption, assuming that all arrival times are at time 0.

Remark. The above theorem leaves open the problem of finding a pseudo-polynomial time algorithm for the problem. It is also open to obtain an approximation algorithm for this case of the problem.

We come now to one of our main results, which utilizes a rather sophisticated algorithm and analysis in order to show that, in contrast with the case of identical arrival times, the LBP with identical deadlines is polynomially solvable. The remainder of this section is devoted to proving the following theorem:

Theorem 15. *The LBP with preemption, under constraint II (one can only work on jobs that can be completed) and objective (3) (go home as early as possible), is solvable in polynomial time if all jobs have the same deadlines ($d_i = D$, for all i).*

We begin with a definition a "forced gap:" There is a *forced gap* starting at time τ if τ is the earliest time such that the total work arriving by time τ is less than τ. This (first) forced gap ends at the arrival time, τ', of the next job. Subsequently, there may be more forced gaps, each determined by considering the scheduling problem that starts at the end, τ', of the previous forced gap. We note that a forced gap can have length zero.

Under the "go home early" objective, we can assume, without loss of generality, that there are no forced gaps, since our problem really begins only at the time τ' that the *last* forced gap ends. (The bureaucrat is certainly not allowed to go home before the end τ' of the last forced gap, since more jobs arrive after τ' that can be processed before their deadlines.) While an optimal schedule may contain gaps that are not forced, the next lemma implies that there exists an optimal schedule having no unforced gaps.

Lemma 2. *Consider the LBP of Theorem 15, and assume that there are no forced gaps. If there is a schedule having makespan T, then there is a schedule with no gaps, also having makespan T.*

Lemma 3. *Consider an LBP of Theorem 15 in which there are no forced gaps. Any feasible schedule can be rearranged so that all completed jobs are ordered by their arrival times, and all incomplete jobs are ordered by their arrival times.*

Proof. The proof uses a simple exchange argument as in the standard proof of optimality for the EDD (Earliest Due Date) policy in traditional scheduling problems. □

Our algorithm checks if there exists a schedule having no gaps that completes exactly at time T. Assume that the jobs $1, \ldots, n$ are labeled so that $a_1 \leq a_2 \cdots \leq a_n$. The main steps of the algorithm are as follows:

The Algorithm:

1. Determine the forced gaps. This allows us to reduce to a problem having no forced gaps, which starts at the end of the last forced gap.
   The forced gaps are readily determined by computing the partial sums, $\tau_j = \sum_{i=1}^{j} t_i$, for $j = 0, 1, \ldots, n$, and comparing them to the arrival times. (We define $\tau_0 = 0$.) The first forced gap, then, begins at the time $\tau = \tau_{j^*} = \min\{\tau_j : \tau_j < a_{j+1}\}$ and ends at time a_{j^*+1}. ($\tau = 0$ if $a_1 > 0$; $\tau = \infty$ if there are no forced gaps.) Subsequent forced gaps, if any, are computed similarly, just by re-zeroing time at τ', and proceeding as with the first forced gap.
2. Let $x = D - T$ be the length of time between the common deadline D and our target makespan T. A job i for which $t_i \leq x$ is called *short*; jobs for which $t_i > x$ are called *long*.
   If it is *not* possible to schedule the set of short jobs, so that each is completed and they are all done by time T, then our algorithm stops and returns "NO," concluding that going home by time T is impossible. Otherwise, we continue with the next step of the algorithm.
   The rationale for this step is the observation that any job of length at most x must be completed in any schedule that permits the bureaucrat to go home by time T, since its critical time occurs at or after time T.
3. Create a schedule $\mathcal{S}$ of all of the jobs, ordered by their arrival times, in which the amount of time spent on job i is t_i if the job is short (so it is done completely) and is $t_i - x$ if the job is long.
   For a long job i, $t_i - x$ is the maximum amount of time that can be spent on this job without committing the bureaucrat to completing the job, i.e., without causing the adjusted critical time of the job to occur after time T.
   If this schedule $\mathcal{S}$ has no gaps and ends at a time after T, then our algorithm stops and returns "YES." A feasible schedule that allows the bureaucrat to go home by time T is readily constructed by "squishing" the schedule that we just constructed: We reduce the amount of time spent on the long jobs, starting with the latest long jobs and working backwards in time, until

the completion time of the last short job exactly equals T. This schedule completes all short jobs (as it should), and does partial work on long jobs, leaving all of them with adjusted critical times that fall *before* time T (and are therefore not possible to resume at time T, so they can be avoided).

4. If the above schedule $\mathcal{S}$ has gaps or ends before time T, then $\mathcal{S}$ is not a feasible schedule for the lazy bureaucrat, so we must continue the algorithm, in order to decide *which* long jobs to complete, if it is possible to go home by time T. We use a dynamic programming algorithm *Schedule-by-T*, which we describe in detail below.

Procedure Schedule-by-T. Let G_i be the sum of the gap lengths that occur before time a_i in schedule $\mathcal{S}$. Then, we know that in order to construct a gapless schedule, at least $\lceil G_i/x \rceil$ long jobs (in addition to the short jobs) from $1, \ldots, i-1$ must be completed. For each i we have such a constraint; collectively, we call these the *gap constraints*.

Claim. If for each gap in schedule $\mathcal{S}$, there are enough long jobs to be completed in order to fill the gap, then a feasible schedule ending at T exists.

We devise a dynamic programming algorithm as follows. Let $T(m, k)$ be the earliest completion time of a schedule that satisfies the following:

(1) It completes by time T;
(2) It uses jobs from the set $\{1, \ldots, k\}$;
(3) It completes exactly m jobs and does no other work (so it may have gaps, making it an infeasible schedule);
(4) It satisfies the gap constraints; and
(5) It completes all short jobs (of size $\leq x$).

The boundary conditions on $T(m, k)$ are given by:

$T(0, 0) = 0$;
$T(0, n) = \infty$, which implies that at least one of the jobs must be completed;
$T(m, 0) = \infty$ for $m > 0$;
$T(m, k) = \infty$ if there exist constraints such that at least $m+1$ jobs from $1, \ldots, k$ must be completed, some of the jobs from $1, \ldots, k$ must be completed because they are short, and some additional jobs may need to be completed because of the gap constraints. Note that this implies that $T(0, k)$ is equal to zero or infinity, depending on whether gap constraints are disobeyed.

In general, $T(m, k)$ is given by selecting the better of two options: $T(m, k) = \min\{\alpha, \beta\}$, where α is the earliest completion time if we choose not to execute job k (which is a legal option only if job k is long), giving

$$\alpha = \begin{cases} T(m, k-1) & \text{if } t_k > x \\ \infty & \text{otherwise,} \end{cases}$$

and β is the earliest completion time if we choose to execute job k (which is a legal option only if the resulting completion time is by time T), giving

$$\beta = \begin{cases} \max(a_k + t_k, T(m-1, k-1) + t_k) & \text{if this quantity is } \leq T \\ \infty & \text{otherwise.} \end{cases}$$

Lemma 4. *There exists a feasible schedule completing at time T if and only if there exists an m for which $T(m,n) < \infty$.*

Proof. If $T(m,n) = \infty$ for all m, then, since the gap constraints apply to any feasible schedule, and it is not possible to find such a schedule for any number of jobs m, there is no feasible schedule that completes on or before T.

If there exists an m for which $T(m,n) < \infty$, let m^* be the smallest such m. Then, by definition, $T(m^*,n) \leq T$. We show that the schedule $\mathcal{S}^*$ obtained by the dynamic program can be made into a feasible schedule ending at T. Consider jobs that are not completed in the schedule $\mathcal{S}^*$; we wish to use some of them to "fill in" the schedule to make it feasible, as follows.

Ordered by arrival times of incomplete jobs, and doing up to $t_i - x$ of each incomplete job, fill in the gaps. Note that by the gap constraints, there is enough work to fill in all gaps. There are two things that may make this schedule infeasible: (i) Some jobs are worked on beyond their critical times, and (ii) the last job to be done must be a completed one.

(i). *Fixing the critical time problem:* Consider a job i that is processed at some time, beginning at τ, after its critical time, c_i. We move all completed job pieces that fall between c_i and T to the end of the schedule, lining them up to end at T; then, we do job i from time c_i up until this batch of completed jobs. This is legal because all completed jobs can be pushed to the end of the schedule, and job i cannot complete once it stops processing.

(ii). *Fixing the last job to be a complete one:* Move a "sliver" of the last completed job to just before time T. If this is not possible (because the job would have to be done before it arrives), then it means that we *must* complete one additional job, so we consider $m^* + 1$, and repeat the process. □

This completes the proof of our main theorem, Theorem 15.

References

1. A. Bar-Noy, S. Guha, J. Naor, and B. Schieber. Approximating the throughput of real-time multiple machine scheduling. In *Proc. 31st ACM Symp. Theory of Computing*, 1999.
2. M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the theory of NP-completeness.* W. H. Freeman, San Francisco, 1979.
3. D. Karger, C. Stein, and J. Wein. Scheduling algorithms. In *CRC Handbook of Computer Science*, 1997. To appear.
4. T. Keneally. *Schindler's List.* Touchstone Publishers, New York, 1993.
5. E. Lawler, J. Lenstra, A. Kan, and D. Shmoys. Sequencing and scheduling: Algorithms and complexity. In *Handbooks of Operations Research and Management Science*, volume 4, pages 445–522. Elsevier Science Publishers B.V., 1993.
6. E. L. Lawler and J. M. Moore. A functional equation and its application to resource allocation and sequencing problems. *Management Science*, 16:77–84, 1969.
7. M. Pinedo. *Scheduling: Theory, Algorithms, and Systems.* Prentice Hall, 1995.