

Unrelated Machine Scheduling with Stochastic Processing Times

Martin Skutella

TU Berlin, Institut für Mathematik, MA 5-2
 Straße des 17. Juni 136, 10623 Berlin, Germany, martin.skutella@tu-berlin.de

Maxim Sviridenko

Yahoo! Labs, 111 West 40th Street, 17th Floor New York, NY 10018, USA, sviri@yahoo-inc.com

Marc Uetz

University of Twente, Dept. Applied Mathematics
 P.O. Box 217, 7500 AE Enschede, The Netherlands, m.uetz@utwente.nl

Two important characteristics encountered in many real-world scheduling problems are heterogeneous processors and a certain degree of uncertainty about the sizes of jobs. In this paper we address both, and study for the first time a scheduling problem that combines the classical unrelated machine scheduling model with stochastic processing times of jobs. By means of a novel time-indexed linear programming relaxation, we compute in polynomial time a scheduling policy with performance guarantee $(3 + \Delta)/2 + \varepsilon$ for the stochastic version of the unrelated parallel machine scheduling problem with the weighted sum of completion times objective. Here, $\varepsilon > 0$ is arbitrarily small, and Δ is an upper bound on the squared coefficient of variation of the processing times. When jobs also have individual release dates, our bound is $(2 + \Delta) + \varepsilon$. We also show that the dependence of the performance guarantees on Δ is tight. Via $\Delta = 0$, the currently best known bounds for deterministic scheduling on unrelated machines are contained as special case.

Key words: stochastic scheduling; approximation algorithm; linear programming relaxation

MSC2000 subject classification: 90B36; 68M20; 90C27; 90C59; 68W25; 68W40; 68Q25

OR/MS subject classification: Primary: production/scheduling ; secondary: approximations/heuristic

History:

1. Introduction The problem to minimize the total weighted completion time on unrelated parallel machines, denoted $R|(r_{ij})|\sum w_j C_j$ in the three-field notation of Graham et al. [8], is one of the most important classical problems in the theory of deterministic scheduling. Each job j has a weight w_j , possibly an individual release date r_{ij} before which job j must not be scheduled on machine i , and the processing time of job j on machine i is p_{ij} . Each job has to be processed nonpreemptively on any one of the machines, and each machine can process at most one job at a time. The objective is to find a schedule minimizing the total weighted completion time $\sum_j w_j C_j$, where C_j denotes the completion time of job j in the schedule. The special case with identical parallel machines is already known to be strongly NP-hard (Lenstra et al. [13]) but there do exist polynomial time approximation schemes [1, 33]. The general setting of unrelated parallel machines turns out to be significantly harder and there is a complexity gap compared to identical parallel machines: Hoogeveen et al. [11] prove MaxSNP-hardness and hence there is no polynomial time approximation scheme. On the positive side, the currently best known approximation algorithms for deterministic unrelated parallel machines have performance guarantees $3/2$ and 2 , for the problem without and with release dates, respectively [4, 26, 28, 30]. Improving these bounds is considered to be among the most important open problems in scheduling (see Schuurman and Woeginger [27]) which is also an indication of the high significance of unrelated parallel machine scheduling.

Stochastic scheduling. We consider for the first time the stochastic variant of unrelated parallel machine scheduling. Here, the processing time of a job j on machine i is given by random variable P_{ij} . In stochastic scheduling, we are asked to compute a non-anticipatory scheduling policy. Intuitively, a non-anticipatory scheduling policy must make its scheduling decisions at time t based on the observed past up to time t as well as the a priori knowledge of the input data of the problem. A policy, however, is not allowed to anticipate information about the future, i.e., the actual realizations of the processing times of jobs which have not yet been completed by time t . For a thorough discussion and definition of non-anticipatory scheduling policies, see subsequent Section 2. As all previous work in the area, we assume that the random variables P_{ij} are stochastically independent across jobs. For any given non-anticipatory scheduling policy, the possible outcome of the objective function $\sum_j w_j C_j$ is a random variable, and our goal is to minimize its expected value, which by linearity of expectation equals $\sum_j w_j \mathbb{E}[C_j]$.

Related Work. Generalizing a well known result of Smith [34] for deterministic single machine scheduling, Rothkopf [22] proved in 1966 that the WSEPT rule (weighted shortest expected processing time first: schedule jobs in order of non-increasing ratios $w_j/\mathbb{E}[P_j]$) minimizes the expected total weighted completion time on a single machine. Apart from Weiss’ results on the asymptotic optimality of WSEPT in stochastic scheduling on identical parallel machines [37, 38], the first constant factor approximation algorithms for stochastic scheduling on identical parallel machines have been obtained in 1999 by Möhring et al. [18]. Next to a linear programming (LP) based analysis of the WSEPT rule, they define list scheduling policies which are based on linear programming relaxations in completion time variables. The performance bounds are constant whenever the coefficients of variation of the jobs’ processing times are bounded by a constant. As usual in stochastic scheduling, all bounds hold with respect to any non-anticipatory scheduling policy. By using an idea from Chekuri et al. [2], that approach was extended to stochastic scheduling problems with precedence constraints by Skutella and Uetz [32]. Subsequently, in line with earlier work by Chou et al. [3], Megow et al. [15] combined the stochastic scheduling model with online scheduling, and derived combinatorial, constant competitive algorithms that are not guided by linear programming relaxations. Yet all results, including the analysis by Megow et al. [15], are based on one and the same linear programming relaxation, namely that of [18]. With respect to the underlying relaxation, Schulz [25] goes one step further, and uses the mean busy time relaxation that was previously used also by Correa and Wagner [5], yet its validity in stochastic scheduling still relies on the validity of the completion time relaxation of [18]. Nevertheless, in comparison to [15], Schulz obtains improved and simpler results through the clever use of an optimal solution to an equivalent time-indexed LP relaxation for deterministic scheduling.

Two other research directions are related to our work, yet for different models and independent of the techniques of [18] as well as ours. One is approximation algorithms for preemptive stochastic scheduling by Megow and Vredeveld [16]. They use a single machine relaxation that is optimally solved by a Gittins index policy, and thereby achieve a competitive ratio of 2 for preemptive online stochastic scheduling on parallel identical machines. The other is work by Scharbrodt et al. [23] and Souza and Steger [35], who analyze the expected competitive ratio rather than the expected performance of a policy Π . In that model, one analyzes the ratio $\mathbb{E}[v(\Pi)/v(\text{Offline-Opt})]$, where $v(\Pi)$ is the objective function value of a policy Π . In this paper, however, we follow [15, 18, 25, 32] and focus on the ratio $\mathbb{E}[v(\Pi)]/\mathbb{E}[v(\Pi^{\text{Opt}})]$ instead, where Π^{Opt} is an optimal, non-anticipatory scheduling policy, and not the point-wise optimal offline solution as in [23, 35]. We refer to [23] for a discussion of the pro’s of their model as opposed to the one considered here. Note that we here restrict to a weaker adversary, since the adversary is also bound to be non-anticipatory. For a thorough discussion of this issue in the context of competitive analysis, see also Koutsoupias and Papadimitriou [12].

stochastic scheduling model	worst case performance guarantee arbitrary P_{ij}	reference $\mathbb{C}\mathbb{V}[P_{ij}] \leq 1$	reference
$\text{P} \mid \mathbb{E}[\sum w_j C_j]$	$1 + \frac{(m-1)(\Delta+1)}{2m}$	$2 - 1/m$	[18]
$\text{P} \mid r_j \mid \mathbb{E}[\sum w_j C_j]$	$2 + \Delta$	3	[25]
$\text{R} \mid \mathbb{E}[\sum w_j C_j]$	$1 + \frac{\Delta+1}{2} + \varepsilon$	$2 + \varepsilon$	this paper
$\text{R} \mid r_{ij} \mid \mathbb{E}[\sum w_j C_j]$	$2 + \Delta + \varepsilon$	$3 + \varepsilon$	this paper

TABLE 1. Performance bounds for nonpreemptive stochastic machine scheduling problems. Parameter $\varepsilon > 0$ can be chosen arbitrarily small. Parameter Δ upper bounds the squared coefficient of variation $\mathbb{C}\mathbb{V}^2[P_{ij}] = \text{Var}[P_{ij}]/\mathbb{E}^2[P_{ij}]$ for all P_{ij} . The third column shows the results for $\mathbb{C}\mathbb{V}[P_{ij}] \leq 1$; e.g., uniform, exponential, or Erlang distributions. As usual in stochastic scheduling, these multiplicative bounds hold with respect to the expected performance of any non-anticipatory scheduling policy.

Note that all results discussed so far are restricted to the model with identical parallel machines. Table 1 gives an overview of currently best known performance bounds in nonpreemptive stochastic scheduling with minsum objective, next to the new results obtained in this paper.

With respect to algorithmic ideas and techniques, the evolution of stochastic scheduling has largely benefited in the past from progress being made for the corresponding deterministic scheduling problems. For example, all LP-based approximation results for stochastic scheduling on identical parallel machines outlined above build upon a class of linear programming relaxations in completion time variables that dates back to Wolsey [39] and Queyranne [20] (for single machine scheduling), and was later generalized to identical parallel machines by Schulz [24] and Hall et al. [10] who also presented LP-based approximation algorithms for deterministic scheduling problems.

Our contribution. We obtain the first approximation algorithms for stochastic scheduling on unrelated parallel machines. Despite the fact that the unrelated parallel machine scheduling model is significantly richer than identical parallel machine scheduling, our bounds essentially match all previous performance bounds that have been obtained for the corresponding stochastic scheduling problems on identical parallel machines; see Table 1. We also give a tight lower bound, showing that the dependence of the performance bound on the squared coefficient of variation Δ is unavoidable for the class of policies that we use. For the first time we completely depart from the LP relaxation of Möhring et al. [18], and show how to put a novel, time-indexed linear programming relaxation to work in stochastic machine scheduling. We are optimistic that this approach will inspire further research and prove useful for other stochastic optimization problems in scheduling and related areas.

Time-indexed linear programming relaxations have played a pivotal role in the development of constant factor approximation algorithms for deterministic scheduling on unrelated parallel machines (see Schulz and Skutella [26]). In spite of that, it remained unclear and a major open problem how to come up with a meaningful time-indexed LP relaxation for stochastic scheduling problems (see Megow [14]). Here the main difficulty is that, in contrast to deterministic schedules that can be fully described by time-indexed 0-1-variables, scheduling *policies* feature a considerably richer structure including complex dependencies between the execution of different jobs which cannot be easily described by time-indexed variables.

In Section 3 we show how to overcome this difficulty and present the first time-indexed LP relaxation for stochastic scheduling on unrelated parallel machines. Here, the value of the time-indexed variable x_{ijt} represents the probability of job j being started on machine i at time t . Note that even for simple scheduling policies like the WSEPT rule, determining this probability is non-trivial. The machine capacity constraints say that each machine can process at most one job at a time, and formulating this is rather easy for deterministic unrelated machine scheduling. The situation

is somewhat more complicated in the stochastic setting, and we require a fair amount of information about the exact probability distributions of random variables P_{ij} in order to formulate that constraint.

Notice that, due to the stochastic nature of processing times, even a schedule produced by an optimal policy can lead to the situation where infinitely many variables x_{ijt} may take positive values. This already happens for one machine when considering the start time of the second of two jobs with exponentially distributed processing time. Nonetheless, we show how to overcome this difficulty. Indeed, we show that we can compute an LP-solution in polynomial time that approximates the optimal LP solution with arbitrary precision.

In Section 4 we discuss how to turn a feasible solution to the time-indexed LP relaxation into a simple scheduling policy. Our approach is inspired by the randomized rounding algorithm for deterministic scheduling on unrelated parallel machines in [26]. Each job j is randomly assigned to a machine i with probability $\sum_t x_{ijt}$; then, on each machine i , the WSEPT policy is used to schedule the jobs assigned to i . The analysis, however, is based on a somewhat more elaborate, random sequencing of jobs which is determined by a two-stage random process. We show how to extend our results to the setting with release dates in Section 5.

The scheduling policies that we use to obtain our results fall into the special class of *fixed assignment policies*. That means that, already at time zero, each job is immediately and irrevocably assigned to a machine. In particular, these policies ignore the additional information that evolves over time in the form of the actual realizations of processing times. Not surprisingly, this ignorance comes at a price. In Section 6 we prove a lower bound of $\Delta/2$ on the performance guarantee of *any* fixed assignment policy. Moreover, we can also show that the LP relaxation that we use can have an optimality gap in the same order of magnitude.

In order to keep the presentation as simple as possible, we ignore release dates and restrict to the problem $R \mid \mathbb{E}[\sum w_j C_j]$ throughout most of the paper. Only in Section 5 we show how release dates can be taken care of in our approach.

In Section 7 we discuss how to execute our policy with polynomial running time while in Section 8 we discuss an alternative linear programming relaxation.

Parallel to Stochastic Knapsack. There is an interesting parallel of our work with that on stochastic knapsack problems. Indeed, stochastic knapsack problems can be reinterpreted as single machine stochastic scheduling problems where all jobs have due date 1, and with weighted earliness objective. The first study of approximation algorithms for stochastic knapsack problems is due to Dean et al. [6], presenting constant factor approximation algorithms along with an analysis of the adaptivity gap. In stochastic scheduling, the adaptivity gap would correspond to the gap between the best static list scheduling policy and an optimal adaptive scheduling policy. Their results are based on a linear programming relaxation that is essentially the deterministic knapsack LP where item sizes and weights are replaced by expected values. In that sense, methodology-wise their linear program parallels that of [18] in stochastic scheduling on parallel machines. Recently, Gupta et al. [9] were able to obtain constant factor approximation algorithms for a much broader class of stochastic knapsack problems (and other problems, too). Key to these results is a time-indexed linear programming relaxation, based on the same type of variables as we use here. It is interesting to note that in their paper as well as in ours, moving from “natural yet simple” LP relaxations to richer time-indexed LP relaxations is key to more general results.

2. Notation and preliminaries We are given a set of jobs J of cardinality n with job weights $w_j \in \mathbb{Z}_{>0}$, $j \in J$, and a set of unrelated parallel machines M of cardinality m . Moreover, we are given a random variable P_{ij} for every job $j \in J$ and every machine $i \in M$. Each job j needs to be executed on any one of the machines $i \in M$, and each machine can process at most one job at a time. If job j is processed on machine i , its processing time is random according to P_{ij} . In some

applications job $j \in J$ cannot be processed on a certain machine $i \in M$, i. e., $\mathbb{E}[P_{ij}] = \infty$. For the sake of simplicity of presentation, we assume in this paper that $\mathbb{E}[P_{ij}]$ is finite for all $i \in M$ and $j \in J$. Yet all presented results also hold for the more general case where $\mathbb{E}[P_{ij}] = \infty$ for certain pairs (i, j) . Later in Section 5, we consider a slightly more general model where each job $j \in J$ can start on machine $i \in M$ only at time $r_{ij} \in \mathbb{Z}_{\geq 0}$ or later. The parameters r_{ij} are called *machine dependent release dates*.

In the stochastic scheduling model, the actual realization of the processing time of a job j is only known upon j 's completion and we are looking for a non-anticipatory scheduling policy Π which minimizes the expected total weighted completion time $\mathbb{E}[\sum_j w_j C_j]$, where C_j denotes the completion time of job j . We next define the notion of a scheduling policy.

Scheduling Policies. We here confine ourselves with the intuitive, dynamic view on scheduling policies that puts stochastic scheduling in the framework of stochastic dynamic optimization; see, e.g., Ross [21], and refer to Möhring et al. [17] for the analytic definition of non-anticipatory stochastic scheduling policies as mappings $\Pi: \mathbb{R}^n \rightarrow \mathbb{R}^n$ that map processing times to completion times of jobs. A scheduling policy takes actions at points in time $t \geq 0$, starting at time $t = 0$. An action at a given time t consists of a set of jobs, possibly empty, to be started on a set of idle machines, together with a tentative next decision time $t^* > t$. The next action of the policy is then due at time t^* , or the time of the next job completion, or the time when the next job is released, whatever occurs first. Depending on the action of the policy, the next decision time as well as the state of the schedule at the next decision time is realized according to the probability distributions of the jobs' processing times. A way to think about this dynamic process is as follows. There is an oracle with access to the realized processing times of the jobs, and at any decision time t with jobs J_t in process or started, it computes $t' := \min_{j \in J_t} C_j$. The next decision time is then $t^{\text{next}} := \min\{t', t^*\}$. The information revealed to policy Π is the time t^{next} , along with the set of jobs finished at t^{next} . If there are also release dates r_{ij} , this is taken into account in the obvious way.

In the class of non-anticipatory scheduling policies, the existence of a policy that minimizes the expected performance for the model considered here follows from Theorem 4.2.6 in [17], because of the linearity of the objective function (which is in particular lower semi-continuous). Observe that such a scheduling policy can be represented by a finite, yet generally exponential size decision tree. Finiteness is not clear a priori unless we restrict the set of scheduling policies, for example to being elementary, where jobs are restricted to be started only upon completion times or release dates (or in other words, $t^* = \infty$). But as we argue below, we can restrict essentially w.l.o.g. to discrete processing time distributions, and then the resulting decision tree of any reasonable policy is indeed finite. Also observe that a policy may gain information about remaining processing times of jobs over time, but at any point in time it has only access to distributional information about remaining processing times of unfinished jobs, conditioned on the state of the schedule at time t . That means in particular that, for a given and fixed state at a time t , a scheduling policy cannot choose different actions dependent on the future after time t . Allowing for randomization, however, is possible, meaning that a policy runs a lottery (independent of the future) among all feasible actions at time t .

A brief, concrete example may help: Imagine a job j which has processing time either small (1) or large (M), both with probability $1/2$. For a scheduling policy that starts this job at time t , it can make sense to define a tentative next decision time at $t^* = t + 1$, because then it knows with certainty what the actual processing time of job j is. Using such building blocks, one can even show that an optimal scheduling policy for the setting considered here is generally not work conserving, i. e., there exist examples where an optimal scheduling policy must leave machines deliberately idle even though there are yet unscheduled jobs; see Uetz [36].

Discretization. Throughout this paper we assume that the random variables P_{ij} , $i \in M$, $j \in J$, take positive integral values only. The following lemma states that this assumption costs at most a factor $1 + \varepsilon$ in the objective function value.

LEMMA 1. *For any fixed $\varepsilon > 0$, while only losing a factor $1 + \varepsilon$ in the objective function value, an arbitrary instance can be modified such that the random variables P_{ij} , $i \in M$, $j \in J$, take positive integral values only.*

Proof. If $\mathbb{E}[P_{ij}] = 0$ and $r_{ij} = 0$ for some pair (i, j) , then we can ignore job j since it can be scheduled at no further cost on machine i at time 0. We can thus assume from now on that $\mathbb{E}[P_{ij}] > 0$ or $r_{ij} > 0$ for all pairs (i, j) . By scaling processing times and release dates appropriately, we can make sure that $\mathbb{E}[P_{ij}] \geq \frac{n}{\varepsilon}$ or $r_{ij} \geq \frac{n}{\varepsilon}$ for each pair (i, j) . As a result of this scaling step we know that, for any scheduling policy, $\mathbb{E}[C_j] \geq n/\varepsilon$ for each job $j \in J$. Rounding up all processing times to the nearest positive integer therefore increases the (expected) completion time of any job j by at most $n \leq \varepsilon \mathbb{E}[C_j]$. The overall increase in the objective function is thus bounded by a factor $1 + \varepsilon$. To be more precise, for the rounded instance we can simulate any given policy Π for the original instance, yielding a new scheduling policy that achieves expected completion times of any job j that exceed those of Π by at most $n \leq \varepsilon \mathbb{E}[C_j]$. This is because the rounding will not yield any “new” states that Π wouldn’t know how to handle. Moreover, any scheduling policy for the rounded instance can be translated back to a scheduling policy for the original instance with the same objective value, by adding deliberate idle time where necessary. This shows that the rounding is indeed no loss of generality, except for the additional multiplicative factor $(1 + \varepsilon)$. \square

Input size and further preliminaries. Given that all processing times are integral, we can obviously assume with no further loss of generality that jobs can only be started at integral points in time $t \in \mathbb{Z}_{>0}$. In order to write down an LP relaxation in time-indexed variables, we require a fair amount of information about the exact probability distributions of random variables P_{ij} . More precisely, besides the expectations $\mathbb{E}[P_{ij}]$, we also need the values

$$q_{ijr} := \Pr[P_{ij} \geq r + 1] \quad \text{for } i \in M, j \in J, \text{ and } r \in \mathbb{Z}_{\geq 0}.$$

This, of course, raises questions about the input size of the problem. Here we make the following assumption. In the input we are given, the expected processing time $\mathbb{E}[P_{ij}]$ for each job $j \in J$ and each machine $i \in M$. Moreover, we have access to an oracle which, for any triple (i, j, r) returns q_{ijr} . We emphasize that, in order for our approach to work, it suffices to get these values within some finite precision at the expense of an additional factor $1 + \varepsilon$ in the performance guarantee of our algorithms. More precisely, it suffices to get the values q_{ijr} rounded to multiples of ε/n , which, in particular, can be encoded polynomially in the input size. Notice that such an oracle can be simulated by a polynomial-time Monte Carlo algorithm that can sample from the distribution of the random variables P_{ij} . Having said that, in order to keep the presentation simple we neglect these aspects throughout the paper and assume that we have access to the exact values q_{ijr} .

In the analysis of our algorithm we need the following standard property of the moments of random variable P_{ij} , the proof of which is based on standard results for the n th moment of a random variable, see, e.g. [7, V.6, Lemma 1]. For the sake of completeness, we present the simple proof here.

LEMMA 2. *Let $j \in J$ and $i \in M$. Then,*

$$\sum_{r \in \mathbb{Z}_{\geq 0}} q_{ijr} = \mathbb{E}[P_{ij}] \quad \text{and} \quad \sum_{r \in \mathbb{Z}_{\geq 0}} (r + \frac{1}{2}) q_{ijr} = \frac{1 + \text{CV}[P_{ij}]^2}{2} \mathbb{E}[P_{ij}]^2,$$

where $\text{CV}[P_{ij}]^2 := (\mathbb{E}[P_{ij}^2] - \mathbb{E}[P_{ij}]^2) / \mathbb{E}[P_{ij}]^2$ is the squared coefficient of variation of P_{ij} .

Proof. First,

$$\sum_{r \geq 0} q_{ijr} = \sum_{r \geq 0} \sum_{q \geq r} \Pr[P_{ij} = q + 1] = \sum_{r \geq 0} (r + 1) \Pr[P_{ij} = r + 1] = \mathbb{E}[P_{ij}] .$$

For the second claim,

$$\sum_{r \geq 0} (r + \frac{1}{2}) q_{ijr} = \sum_{r \geq 0} \sum_{q \geq r} (r + \frac{1}{2}) \Pr[P_{ij} = q + 1] = \sum_{r \geq 0} \frac{1}{2} (r + 1)^2 \Pr[P_{ij} = r + 1] = \frac{1}{2} \mathbb{E}[P_{ij}^2] .$$

The claim now follows by definition of the coefficient of variation. \square

3. Time-indexed LP relaxation In the following we derive an LP relaxation of the stochastic scheduling problem under consideration. For a given non-anticipatory scheduling policy Π , let x_{ijt} be the probability that Π starts job $j \in J$ on machine $i \in M$ at time $t \in \mathbb{Z}_{\geq 0}$. Notice that this random decision may depend on the actual processing times of other jobs started by Π before time t . On the other hand, due to the non-anticipatory nature of policy Π , the random variable P_{ij} is independent of Π 's random decision to start job j on machine i at time t .

As the x_{ijt} 's are going to be the variables of our LP relaxation, we derive crucial properties that are going to be the constraints of the LP relaxation. If job $j \in J$ is started on machine $i \in M$ at time $t \in \mathbb{Z}_{\geq 0}$, due to the non-anticipatory nature of policy Π , j 's expected completion time is $t + \mathbb{E}[P_{ij}]$. Thus, by linearity of expectation, the expected completion time of j is

$$\mathbb{E}[C_j] = \sum_{i \in M} \sum_{t \in \mathbb{Z}_{\geq 0}} x_{ijt} (t + \mathbb{E}[P_{ij}]) .$$

A more careful look at j 's behavior reveals the following property. Conditioning on j being started on machine i at time t , the probability that j is still occupying machine i within the later time interval $[s, s + 1]$, $s \in \mathbb{Z}_{\geq t}$, is equal to q_{ijs-t} by definition. Unconditioning yields

$$\Pr[i \text{ processes } j \text{ in } [s, s + 1]] = \sum_{t=0}^s x_{ijt} q_{ijs-t} . \quad (1)$$

As machine i can process at most one job at a time, also the *expected* number of jobs being processed by i in $[s, s + 1]$ is bounded by 1. That is, by linearity of expectation,

$$\sum_{j \in J} \sum_{t=0}^s x_{ijt} q_{ijs-t} \leq 1 .$$

Finally, since policy Π has to process all jobs, we get $\sum_{i \in M} \sum_{t \in \mathbb{Z}_{\geq 0}} x_{ijt} = 1$, for every job j . Thus, the probabilities x_{ijt} corresponding to policy Π form a feasible solution to the following LP relaxation, and the value of this LP solution x is equal to the expected value of the schedule produced by policy Π :

$$\begin{aligned} \min \quad & \sum_{j \in J} w_j C_j^{\text{LP}} \\ \text{s.t.} \quad & \sum_{i \in M} \sum_{t \in \mathbb{Z}_{\geq 0}} x_{ijt} = 1 && \text{for all } j \in J, \end{aligned} \quad (2)$$

$$\sum_{j \in J} \sum_{t=0}^s x_{ijt} q_{ijs-t} \leq 1 \quad \text{for all } i \in M, s \in \mathbb{Z}_{\geq 0}, \quad (3)$$

$$C_j^{\text{LP}} = \sum_{i \in M} \sum_{t \in \mathbb{Z}_{\geq 0}} x_{ijt} (t + \mathbb{E}[P_{ij}]) \quad \text{for all } j \in J, \quad (4)$$

$$x_{ijt} \geq 0 \quad \text{for all } j \in J, i \in M, t \in \mathbb{Z}_{\geq 0}.$$

Notice that the LP variables C_j^{LP} are uniquely determined by the x -variables and could as well be omitted by replacing them in the objective function with the right hand side of (4).

Also notice that this LP is not a formulation but only relaxation of the scheduling problem. This can be seen by realizing that the LP allows to greedily “pack” fractions of jobs in a way that cannot be translated back into any feasible scheduling policy. This is exactly what is exploited in the example that we give in the proof of Theorem 6, where we show that the LP has a large optimality gap even on a single machine. Intuitively, the LP is allowed to neglect implicit temporal dependencies that upper bound the probabilities for jobs being started. For instance in the example where we have two jobs with exponentially distributed processing times that must be scheduled on a single machine, we can define LP solutions where both jobs are started with probability one at some finite point in time. However in any schedule, for any time $t \geq 0$ there is nonzero probability for one job to be started at time t or later.

That said, it becomes clear that the linear program generally suffers from infinitely many variables and constraints. Indeed, it is true that infinitely many variables may be needed in order to map an optimal scheduling policy Π to a corresponding LP solution. Consider again the same two jobs with exponentially distributed processing times that must be scheduled on a single machine, and consider the policy that schedules the jobs one after another. Then the start time of the second job cannot be bounded, i. e., for any $t \geq 0$, the second job is started with positive probability at time t or later. Despite this peculiarity, we can show that an *optimal* solution to the LP relaxation does get along with finite support. More precisely, we give a pseudo-polynomial upper bound on the largest time index t such that $x_{ijt} > 0$ in an optimal LP solution for any $j \in J$ and $i \in M$. Thereby, the problem can be overcome at the expense of an additional factor $1 + \varepsilon$ in the performance guarantee of our algorithms.

THEOREM 1. *The above infinite time-indexed LP relaxation can be solved in pseudo-polynomial time in the input size. Moreover, a $(1 + \varepsilon)$ -approximate LP solution can be found in time polynomial in the input size and $1/\varepsilon$.*

Proof. Any reasonable scheduling policy produces a schedule where the expected completion time of every job is at most

$$D := \max_{i \in M} \sum_{j \in J} \mathbb{E}[P_{ij}] .$$

In particular, the expected value of the schedule is at most $U := D \sum_{j \in J} w_j$ which is thus also an upper bound on the optimal LP solution. Moreover, let

$$R := 2n \max_{j \in J, i \in M} \mathbb{E}[P_{ij}] .$$

LEMMA 3. *Let $T := 2U + R$. There is an optimal solution x to the LP relaxation such that $x_{ijt} = 0$ for $i \in M$, $j \in J$, and $t > T$.*

Proof. Consider an optimal LP solution x and an arbitrary machine $i \in M$. Since U is an upper bound on the value corresponding to x and since $w_j \geq 1$ for each job $j \in J$, it follows from Markov’s inequality that

$$\sum_{j \in J} \sum_{t \geq 2U} x_{ijt} \leq \frac{1}{2} . \quad (5)$$

Using Markov’s inequality once again, we derive that for each job $j \in J$ and $r \geq R$

$$q_{ijr} = \Pr[P_{ij} \geq r + 1] \leq \Pr[P_{ij} \geq 2n \mathbb{E}[P_{ij}]] \leq \frac{1}{2n} . \quad (6)$$

We now define a new LP solution x' by letting

$$x'_{ijt} := \begin{cases} x_{ijt} & \text{for } t < T, \\ \sum_{t' \geq T} x_{ijt'} & \text{for } t = T, \\ 0 & \text{for } t > T. \end{cases}$$

By definition, x' has the desired property stated in the lemma and its objective function value is bounded from above by the value of x . It remains to prove that x' is a feasible LP solution. It is clear that x' fulfills LP constraints (2) since x fulfills these constraints. It is, however, less obvious that x' also fulfills all LP constraints (3). For $s < T$

$$\sum_{j \in J} \sum_{t=0}^s x'_{ijt} q_{ij s-t} = \sum_{j \in J} \sum_{t=0}^s x_{ijt} q_{ij s-t} \leq 1$$

since x is a feasible LP solution and thus satisfies (3). Finally, for $s \geq T$,

$$\begin{aligned} \sum_{j \in J} \sum_{t=0}^s x'_{ijt} q_{ij s-t} &= \sum_{j \in J} \sum_{t=0}^{2U-1} x_{ijt} q_{ij s-t} + \sum_{j \in J} \sum_{t=2U}^s x'_{ijt} q_{ij s-t} \\ &\leq \frac{1}{2n} \sum_{j \in J} \sum_{t=0}^{2U-1} x_{ijt} + \sum_{j \in J} \sum_{t \geq 2U} x_{ijt} \leq \frac{1}{2} + \frac{1}{2} = 1 \end{aligned}$$

where the first inequality follows from (6) and the definition of x' , and the last inequality follows from (2) and (5). This concludes the proof of Lemma 3. \square

It can be easily derived from the proof that the property claimed in Lemma 3 indeed holds for any optimal LP solution x . It is important to notice that, for such a solution x , constraints (3) for $s > T$ are implied by the constraint for $s = T$ since q_{ijr} is non-increasing in r . We have thus reduced the problem of finding an optimal LP solution to solving a *truncated* time-indexed LP of pseudo-polynomial size. It is well known that such time-indexed LPs can be solved approximately in polynomial time at the expense of losing a factor $1 + \varepsilon$ in the objective function. The underlying idea is to replace the discretization of time into unit size intervals by a slightly rougher discretization using intervals of geometrically increasing lengths. We refer to [29, Chapter 2.13] for a more thorough discussion of this point. This concludes the proof of Theorem 1. \square

We finally mention that Theorem 1 can be easily generalized to the more general problem with release dates studied in Section 5.

4. Turning an LP solution into a scheduling policy For a feasible LP solution x , let $X_{ij} := \sum_{t \in \mathbb{Z}_{\geq 0}} x_{ijt}$ for $i \in M, j \in J$. LP constraints (2) imply that $\sum_{i \in M} X_{ij} = 1$ for every job $j \in J$.

Given the values X_{ij} corresponding to a feasible LP solution x , scheduling policy ASSIGN(X) assigns each job $j \in J$ independently at random to one machine $i \in M$ with probability X_{ij} . Then, on each machine $i \in M$, it sequences jobs assigned to i according to the WSEPT rule. In order to formulate our main theorem, remember that Δ upper bounds the squared coefficient of variation $\text{CV}[P_{ij}]^2$ for all P_{ij} .

THEOREM 2. *The expected value of the schedule constructed by policy ASSIGN(X) is at most $\frac{3}{2} + \frac{\Delta}{2}$ times the value of the underlying LP solution x . Thus, by Theorem 1, for any given instance of the stochastic scheduling problem and for any $\varepsilon > 0$, a $(\frac{3}{2} + \frac{\Delta}{2} + \varepsilon)$ -approximate scheduling policy can be found in polynomial time.*

It is not difficult to see that, instead of the random assignment of jobs to machines, we can use a deterministic assignment obtained via the method of conditional probabilities and still get the same performance guarantee. A similar approach for deterministic scheduling was used by Schulz and Skutella [26].

The proof of Theorem 2 is based on a refined, somewhat more complicated policy, that takes the entire LP solution x into account and yields a worse schedule in expectation. It is therefore sufficient to prove the bound stated in Theorem 2 for this alternative policy which we refer to as `ASSIGN&SEQUENCE`(x).

`Assign&Sequence`(x)

1. For every job $j \in J$, choose a pair (i, t) independently at random with probability x_{ijt} and some $r \in \mathbb{Z}_{\geq 0}$ independently at random with probability $q_{ijr}/\mathbb{E}[P_{ij}]$; assign job j to machine i and set its *tentative start time* s to $s := t + r$ (we write “ $j \rightarrow (i, s)$ ” for short).

2. On each machine $i \in M$, sequence all jobs assigned to i in order of increasing tentative start times; ties are broken randomly.

Notice that, as in the simpler policy `ASSIGN`(X), job j is assigned to machine i with probability $\sum_{t \in \mathbb{Z}_{\geq 0}} x_{ijt} = X_{ij}$. Since `ASSIGN`(X) sequences the jobs on every machine in an optimal way, it is superior to policy `ASSIGN&SEQUENCE`(x). By construction of policy `ASSIGN&SEQUENCE`(x), the probability of assigning job $j \in J$ to machine $i \in M$ and setting its tentative start time to $s \in \mathbb{Z}_{\geq 0}$ is

$$\Pr[j \rightarrow (i, s)] = \sum_{t=0}^s x_{ijt} \frac{q_{ijs-t}}{\mathbb{E}[P_{ij}]} . \quad (7)$$

We prove the following job-by-job performance guarantee for `ASSIGN&SEQUENCE`(x).

THEOREM 3. *For every job $j \in J$, the expected value of j 's completion time in the schedule constructed by policy `ASSIGN&SEQUENCE`(x) is at most $(\frac{3}{2} + \frac{\Delta_j}{2}) C_j^{\text{LP}}$ where $\Delta_j := \max_{i \in M} \text{CV}[P_{ij}]^2$.*

By linearity of expectation, Theorem 3 immediately implies Theorem 2. In the proof of Theorem 3 we make use of the following lemma.

LEMMA 4. *Let $j \in J$, $i \in M$, and $s \in \mathbb{Z}_{\geq 0}$. If $j \rightarrow (i, s)$, then the expected total processing time of jobs that policy `ASSIGN&SEQUENCE`(x) schedules on machine i before job j is at most $s + \frac{1}{2}$.*

Proof. We first bound the expected total processing time of jobs $k \neq j$ with $k \rightarrow (i, s')$ for some fixed $s' \in \mathbb{Z}_{\geq 0}$:

$$\sum_{k \neq j} \mathbb{E}[P_{ik}] \Pr[k \rightarrow (i, s')] \stackrel{(7)}{=} \sum_{k \neq j} \sum_{t'=0}^{s'} x_{ikt'} q_{iks'-t'} \leq 1 \quad \text{by (3)}.$$

Thus, the expected total processing times of jobs processed before job j on machine i is at most

$$\sum_{k \neq j} \mathbb{E}[P_{ik}] \left(\sum_{s'=0}^{s-1} \Pr[k \rightarrow (i, s')] + \frac{1}{2} \Pr[k \rightarrow (i, s)] \right) \leq s + \frac{1}{2} ,$$

where the expectation is taken with respect to both the random decisions of policy `ASSIGN&SEQUENCE`(x) as well as the random processing times of jobs $k \neq j$. This concludes the proof of Lemma 4. \square

Proof of Theorem 3. By Lemma 4 we get

$$\mathbb{E}[C_j \mid j \rightarrow (i, s)] \leq s + \frac{1}{2} + \mathbb{E}[P_{ij}] \quad (8)$$

for every job $j \in J$, machine $i \in M$, and tentative start time $s \in \mathbb{Z}_{\geq 0}$. Unconditioning the expectation yields

$$\mathbb{E}[C_j] = \sum_{i \in M} \sum_{s \in \mathbb{Z}_{\geq 0}} \mathbb{E}[C_j \mid j \rightarrow (i, s)] \Pr[j \rightarrow (i, s)] .$$

Applying inequality (8) and equation (7) we get

$$\mathbb{E}[C_j] \leq \sum_{i=1}^m \sum_{s \in \mathbb{Z}_{\geq 0}} \left(s + \frac{1}{2} + \mathbb{E}[P_{ij}] \right) \sum_{t=0}^s x_{ijt} \frac{q_{ijs-t}}{\mathbb{E}[P_{ij}]} .$$

Exchanging the order of summation of s and t , and setting $r := s - t$ yields

$$\begin{aligned} \mathbb{E}[C_j] &\leq \sum_{i=1}^m \sum_{t \in \mathbb{Z}_{\geq 0}} x_{ijt} \left(t + \mathbb{E}[P_{ij}] + \sum_{r \in \mathbb{Z}_{\geq 0}} (r + \frac{1}{2}) \frac{q_{ijr}}{\mathbb{E}[P_{ij}]} \right) \\ &= \sum_{i=1}^m \sum_{t \in \mathbb{Z}_{\geq 0}} x_{ijt} \left(t + \left(\frac{3}{2} + \frac{\mathbb{CV}[P_{ij}]^2}{2} \right) \mathbb{E}[P_{ij}] \right) \\ &\leq \left(\frac{3}{2} + \frac{\Delta_j}{2} \right) C_j^{\text{LP}} \end{aligned}$$

by Lemma 2 and (4). This concludes the proof. \square

We note that the same results can in fact be obtained by considering a weaker LP relaxation in variables y_{ijs} , corresponding to the probability that job j is being processed on machine i in time interval $[s, s + 1]$; see Section 8.

5. Adding release dates In this section we show how to adapt our analysis for a more general problem where each job $j \in J$ comes with a machine dependent deterministic release date $r_{ij} \in \mathbb{Z}_{\geq 0}$ before which job j must not be scheduled on machine i . To handle release dates we add one additional family of constraints to our time-indexed LP relaxation:

$$x_{ijt} = 0 \quad \text{for all } i \in M, j \in J, t < r_{ij}.$$

These constraints are obviously fulfilled by the probabilities x_{ijt} corresponding to an arbitrary scheduling policy Π as no job may be started before it is released. We consider the same LP based policy $\text{ASSIGN\&SEQUENCE}(x)$ for this more general problem.

THEOREM 4. *In the presence of release dates, for every job $j \in J$, the expected value of j 's completion time in the schedule constructed by policy $\text{ASSIGN\&SEQUENCE}(x)$ is at most $(2 + \Delta_j) C_j^{\text{LP}}$ where $\Delta_j := \max_{i \in M} \mathbb{CV}[P_{ij}]^2$.*

The proof of Theorem 4 is almost identical to the proof of Theorem 3; for the sake of completeness we nevertheless present it here.

Proof of Theorem 4. Note that the release dates of all jobs that have tentative start times less than s is at most s . Thus, by Lemma 4, we get

$$\mathbb{E}[C_j \mid j \rightarrow (i, s)] \leq s + s + \frac{1}{2} + \mathbb{E}[P_{ij}] \quad (9)$$

for every job $j \in J$, machine $i \in M$, and tentative start time $s \in \mathbb{Z}_{\geq 0}$. Unconditioning the expectation yields

$$\mathbb{E}[C_j] = \sum_{i \in M} \sum_{s \in \mathbb{Z}_{\geq 0}} \mathbb{E}[C_j | j \rightarrow (i, s)] \Pr[j \rightarrow (i, s)] .$$

Applying inequality (9) and equation (7) we get

$$\mathbb{E}[C_j] \leq \sum_{i=1}^m \sum_{s \in \mathbb{Z}_{\geq 0}} (2s + \frac{1}{2} + \mathbb{E}[P_{ij}]) \sum_{t=0}^s x_{ijt} \frac{q_{ij s-t}}{\mathbb{E}[P_{ij}]} .$$

Exchanging the order of summation of s and t , and setting $r := s - t$ yields

$$\begin{aligned} \mathbb{E}[C_j] &\leq \sum_{i=1}^m \sum_{t \in \mathbb{Z}_{\geq 0}} x_{ijt} \left(2t + \mathbb{E}[P_{ij}] + 2 \sum_{r \in \mathbb{Z}_{\geq 0}} (r + \frac{1}{2}) \frac{q_{ijr}}{\mathbb{E}[P_{ij}]} \right) \\ &= \sum_{i=1}^m \sum_{t \in \mathbb{Z}_{\geq 0}} x_{ijt} \left(2t + (2 + \mathbb{CV}[P_{ij}]^2) \mathbb{E}[P_{ij}] \right) \leq (2 + \Delta_j) C_j^{\text{LP}} \end{aligned}$$

by Lemma 2 and (4). This concludes the proof of Theorem 4. \square

We conclude this section with the following result for the model with release dates.

THEOREM 5. *In the presence of release dates, the expected value of the schedule constructed by policy ASSIGN&SEQUENCE(x) is at most $2 + \Delta$ times the value of the underlying LP solution x . Thus, for any given instance of the stochastic scheduling problem and for any $\varepsilon > 0$, a $(2 + \Delta + \varepsilon)$ -approximate scheduling policy can be found in polynomial time.*

6. Tightness of Performance Bounds We argue that our results cannot be easily improved, because both LP relaxation as well as our scheduling policies have an optimality gap of $\Theta(\Delta)$. The following theorem is somewhat surprising since the corresponding time-indexed linear program for the deterministic single machine scheduling problem has the same optimal value as an optimal schedule.

THEOREM 6. *Even for the special case of a single machine, the multiplicative gap between the expected value of an optimal policy and the value of an optimal LP solution can be as large as $\Delta/2$.*

Proof of Theorem 6. Consider the following single machine instance. We are given a set of n identical jobs $J = \{0, \dots, n-1\}$ with unit weight $w_j = 1$ and stochastic processing times

$$P_j = \begin{cases} 1 & \text{with probability } 1 - 1/(n+1), \\ n^4 + n^3 - n & \text{with probability } 1/(n+1), \end{cases} \quad \text{for all } j \in J.$$

In particular, $\mathbb{E}[P_j] = n^3$, $\mathbb{E}[P_j^2] = n^7 + n^6 - 2n^4 + n$, and $\mathbb{CV}[P_j]^2 = n - 2/n^2 + 1/n^5$. Also notice that

$$q_{jr} = \begin{cases} 1 & \text{for } r = 0, \\ 1/(n+1) & \text{for } r = 1, \dots, n^4 + n^3 - n - 1, \\ 0 & \text{otherwise,} \end{cases} \quad \text{for all } j \in J.$$

Moreover, the expected value of the schedule found by the optimal WSEPT policy is equal to

$$\sum_{j=0}^{n-1} (j+1) n^3 = \frac{1}{2} n^4 (n+1) .$$

Consider the LP solution given by

$$x_{j,t} = \begin{cases} 1/n & \text{for } t = jn, \dots, (j+1)n - 1, \\ 0 & \text{otherwise,} \end{cases} \quad \text{for all } j \in J.$$

It is not difficult to check that x is a feasible LP solution: Notice that LP constraints (2) are fulfilled by definition of x . Moreover, the left-hand side of LP constraints (3) is maximal for $s = n^2 - 1$ where it takes the value

$$\sum_{j=0}^{n-1} \sum_{t=0}^{n^2-1} x_{j,t} q_{j,n^2-1-t} = \sum_{j=0}^{n-2} \frac{1}{n+1} + \sum_{t=0}^{n^2-2} \frac{x_{n-1,t}}{n+1} + x_{n-1,n^2-1} = \frac{n-1}{n+1} + \frac{n-1}{n(n+1)} + \frac{1}{n} = 1 .$$

The value of the LP solution x is

$$\sum_{j=0}^{n-1} \sum_{t=jn}^{(j+1)n-1} \frac{1}{n} (t + n^3) = \frac{1}{n} \sum_{t=0}^{n^2-1} (t + n^3) = \frac{n^2(n^2-1)}{2n} + n^4 = n^4 + \frac{1}{2}n^3 - \frac{1}{2}n .$$

Thus, the multiplicative gap between the expected value of the optimal WSEPT policy and the value of an optimal LP solution is at least

$$\frac{n^4(n+1)}{2n^4+n^3-n} = \frac{n}{2} + \Theta(1) = \frac{\Delta}{2} + \Theta(1) ,$$

where $\Delta = \mathbb{C}\mathbb{V}[P_j]^2 = n - 2/n^2 + 1/n^5$. \square

Next we show that the dependence of our performance guarantee on the squared coefficients of variation $\Delta = \max_{ij} \mathbb{C}\mathbb{V}[P_{ij}]^2$ has the right order of magnitude for all algorithms that use fixed job assignments to machines, i. e., for all policies where the assignment of all jobs to machines is fixed right in the beginning and not changed after learning about the processing time realizations of already completed jobs. This even holds for the case of identical parallel machines where, for each job $j \in J$, there is one random variable P_j such that $P_{ij} = P_j$ for all machines $i \in M$. The following theorem shows that our approximation result cannot be significantly improved without considering adaptive policies.

THEOREM 7. *Even for the special case of identical parallel machines, the performance ratio of any fixed-assignment policy can be as large as $\frac{(1-\delta)\Delta}{2}$ for any $\delta > 0$, for large enough number of machines m .*

Proof of Theorem 7. Let $\varepsilon > 0$ be a fixed small constant. Consider an instance consisting of m identical parallel machines and m^2 jobs with unit weights $w_j = 1$ and stochastic processing times P_j with

$$P_j = \begin{cases} 1 & \text{with probability } \frac{1-\varepsilon}{m}, \\ 0 & \text{with probability } 1 - \frac{1-\varepsilon}{m}. \end{cases}$$

The expected processing time of each job is $\mathbb{E}[P_j] = \Pr[P_j = 1] = \frac{1-\varepsilon}{m}$ and the total expected processing time of all jobs is

$$\mathbb{E} \left[\sum_{j \in J} P_j \right] = (1 - \varepsilon) m .$$

The squared coefficient of variation of random variable P_j is

$$\Delta = \mathbb{C}\mathbb{V}[P_j]^2 = \frac{m}{1-\varepsilon} - 1 .$$

The probability that our instance has at least $\frac{1}{1-\varepsilon} \mathbb{E}[\sum_{j \in J} P_j] = m$ jobs with processing times equal to one is upper bounded by $e^{-\frac{1}{4}m\varepsilon^2/(1-\varepsilon)}$ by the Chernoff bounds for the sum of independent Boolean random variables (see, e. g., [19, Theorems 4.1 and 4.3]). Note also that, under any realization of processing times, the value of the schedule computed by the optimal policy is upper bounded by m^3 .

In order to derive an upper bound on the value of an optimal policy, consider the following adaptive (but non-anticipatory) policy. Initially all machines are available. Start one job on each available machine. If some of the jobs are immediately finished (i. e., they have processing time 0), then start one job per available machine again. Once we have a job that has not been finished immediately, we declare the machine where this job is processed *unavailable* and continue to assign jobs to available machines. The expected value of the schedule produced by the optimal policy can thus be bounded from above by

$$\begin{aligned} & \Pr \left[\sum_{j \in J} P_j < m \right] \mathbb{E} \left[\sum_{j \in J} P_j \mid \sum_{j \in J} P_j < m \right] + \Pr \left[\sum_{j \in J} P_j \geq m \right] m^3 \\ & \leq \mathbb{E} \left[\sum_{j \in J} P_j \right] + \Pr \left[\sum_{j \in J} P_j \geq m \right] m^3 \\ & \leq (1 - \varepsilon) m + e^{-\frac{1}{4}m\varepsilon^2/(1-\varepsilon)} m^3 < m , \end{aligned}$$

where the last inequality holds for large enough m .

Consider now any fixed-assignment policy that assigns k_i jobs to be processed on machine $i \in M$ with $\sum_{i \in M} k_i = m^2$. Since all jobs have identical distributions and weights, the optimal single machine policy is to process jobs according to an arbitrary permutation. The expected value of such a schedule on machine $i \in M$ is

$$\sum_{q=1}^{k_i} q \frac{1-\varepsilon}{m} = (1-\varepsilon) \frac{k_i(k_i+1)}{2m} .$$

Thus, due to convexity, the value of the fixed assignment policy is

$$(1-\varepsilon) \sum_{i \in M} \frac{k_i(k_i+1)}{2m} \geq (1-\varepsilon) \frac{m(m+1)}{2} .$$

Therefore, the worst-case ratio between an optimal fixed assignment policy and an optimal policy is at least $(1-\varepsilon) \frac{m+1}{2}$. Since we can choose $\varepsilon > 0$ to be arbitrarily small and m arbitrarily large, we derive that this ratio is at least $\frac{(1-\delta)\Delta}{2}$ for any $\delta > 0$ and large enough parameter m . \square

7. Execution of Scheduling Policies We have argued that the policy we propose can be computed in polynomial time, but so far did not discuss the computation time to actually execute the scheduling policy, or more generally, any stochastic scheduling policy. The major issue is how, and with which computational effort the scheduler learns about the next job completion when executing a set of jobs. Probabilistically, this event is described by the minimum of a set of random variables, of which we just sample while executing the policy. In general, and already if there is just one single job to be processed, there might of course be nonzero probability for a job to be exponentially longer than expected. But due to Markov's inequality, the probability for exceeding the expected processing time by an exponential factor is exponentially small, too. Therefore, with high probability the sampled processing times of jobs can be encoded polynomially in the input size of the problem. Apart from this minor issue inherent in all stochastic scheduling problems, we note that the policy $\text{ASSIGN}(X)$ is in particular elementary [17], meaning that jobs are only started upon release times or completion times of other jobs. Hence, there is only a linear number of decision times.

8. A weaker time-indexed LP relaxation In this section we discuss a weaker time-indexed LP relaxation of the stochastic scheduling problem. Instead of using variables x_{ijt} corresponding to the probability that job j is started on machine i at time t as in Section 3, we use variables y_{ijs} corresponding to the probability that job j is being processed on machine i within time interval $[s, s + 1]$. Then, by (1), we get

$$y_{ijs} = \sum_{t=0}^s x_{ijt} q_{ij\ s-t} . \quad (10)$$

Thus, the machine capacity constraints (3) can now simply be written as

$$\sum_{j \in J} y_{ijs} \leq 1 \quad \text{for all } i \in M, s \in \mathbb{Z}_{\geq 0}. \quad (11)$$

Moreover, making use of the first part of Lemma 2, constraints (2) translate to

$$\sum_{i \in M} \sum_{s \in \mathbb{Z}_{\geq 0}} \frac{y_{ijs}}{\mathbb{E}[P_{ij}]} = 1 \quad \text{for all } j \in J. \quad (12)$$

Finally, with the help of the second part of Lemma 2, the expected completion time (4) of job j can be rewritten as

$$C_j^{\text{LP}} = \sum_{i \in M} \sum_{s \in \mathbb{Z}_{\geq 0}} \left(\frac{y_{ijs}}{\mathbb{E}[P_{ij}]} \left(s + \frac{1}{2} \right) + \frac{1 - \text{CV}[P_{ij}]^2}{2} y_{ijs} \right) \quad \text{for all } j \in J. \quad (13)$$

Thus, we get the following new LP relaxation:

$$\begin{aligned} \min \quad & \sum_{j \in J} w_j C_j^{\text{LP}} \\ \text{s.t.} \quad & (11), (12), (13) \\ & y_{ijs} \geq 0 \quad \text{for all } j \in J, i \in M, s \in \mathbb{Z}_{\geq 0}. \end{aligned}$$

As any feasible solution x to the LP relaxation from Section 3 can be mapped to a feasible solution y of the same value via (10), the new relaxation is not stronger than the old one. Moreover, there exist deterministic instances ($\Delta = 0$) on identical parallel machines for which the gap between the optimal solution values of the two relaxations is arbitrarily close to 2. For example, Consider one job j with weight $w_j = 1$, processing time $p_j = m$, and m identical parallel machines. The optimal LP solution of the linear program in x -variables has value m which is also the value of an optimal schedule. The optimal value of the linear program in y -variables, however, is $\frac{m+1}{2}$. The new LP can, however, be strengthened by adding the following constraints

$$C_j^{\text{LP}} \geq \sum_{i \in M} \sum_{s \in \mathbb{Z}_{\geq 0}} y_{ijs} \quad \text{for all } j \in J. \quad (14)$$

The algorithms presented in Sections 4 and 5 can be easily reinterpreted with respect to a solution y to the new LP relaxation and yield the same performance bounds.

We finally also remark that due to our insights about the LP relaxation in y -variables, the work of Skutella [30] on approximation algorithms based on convex quadratic programming relaxations can be generalized to the setting of stochastic unrelated machine scheduling.

Acknowledgments. The authors would like to thank Nicole Megow for helpful discussions on the topic of this paper, and Andreas S. Schulz and Benjamin Labonté for valuable comments on an earlier version of this paper. Most of the results presented in this paper were obtained during the Dagstuhl Seminar 13111 “Scheduling” in 2013. The authors would like to thank the organizers and Schloss Dagstuhl for providing an enjoyable and stimulating atmosphere.

Martin Skutella was supported by the DFG Research Center MATHEON “Mathematics for key technologies” in Berlin and by the DFG Focus Program 1307 within the project “Algorithm Engineering for Realtime Scheduling and Routing”. Maxim Sviridenko was supported by EPSRC grants EP/J021814/1, EP/D063191/1, FP7 Marie Curie Career Integration Grant and Royal Society Wolfson Research Merit Award. Marc Uetz was supported by the Centre for Telematics and Information Technology (CTIT) of the University of Twente.

An extended abstract with parts of this paper appeared in the Proc. of the 31st Symposium on Theoretical Aspects of Computer Science, STACS 2014 [31].

References

- [1] Afrati, F., E. Bampis, C. Chekuri, D. Karger, C. Kenyon, S. Khanna, I. Milis, M. Queyranne, M. Skutella, C. Stein, M. Sviridenko. 1999. Approximation schemes for minimizing average weighted completion time with release dates. *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science*. New York City, NY, 32–43.
- [2] Chekuri, C., R. Motwani, B. Natarajan, C. Stein. 2001. Approximation techniques for average completion time scheduling. *SIAM Journal on Computing* **31** 146–166.
- [3] Chou, C.F., H. Liu, M. Queyranne, D. Simchi-Levi. 2006. On the asymptotic optimality of a simple on-line algorithm for the stochastic single machine weighted completion time problem and its extensions. *Operations Research* **54** 464–474.
- [4] Chudak, F. A. 1999. A min-sum $3/2$ -approximation algorithm for scheduling unrelated parallel machines. *Journal of Scheduling* **2** 73–77.
- [5] Correa, J., M. Wagner. 2008. LP-based online scheduling: From single to parallel machines. *Mathematical Programming* **119** 109–136.
- [6] Dean, B., M. X. Goemans, J. Vondrák. 2008. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Mathematics of Operations Research* **33** 945–964.
- [7] Feller, W. 1971. *An Introduction to Probability Theory and Its Applications*, vol. 2. 2nd ed. Wiley Series in Probability and Mathematical Statistics.
- [8] Graham, R. L., E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan. 1979. Optimization and approximation in deterministic sequencing and scheduling: A survey. *Annals of Discrete Mathematics* **5** 287–326.
- [9] Gupta, A., R. Krishnaswamy, M. Molinaro, R. Ravi. 2011. Approximation algorithms for correlated knapsacks and non-martingale bandits. *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society, 827–836.
- [10] Hall, L. A., A. S. Schulz, D. B. Shmoys, J. Wein. 1997. Scheduling to minimize average completion time: Off-line and on-line approximation algorithms. *Mathematics of Operations Research* **22** 513–544.
- [11] Hoogeveen, H., P. Schuurman, G. J. Woeginger. 2001. Non-approximability results for scheduling problems with minsum criteria. *INFORMS Journal on Computing* **13** 157–168.
- [12] Koutsoupias, E., C. H. Papadimitriou. 2000. Beyond competitive analysis. *SIAM Journal on Computing* **30** 300–317.
- [13] Lenstra, J. K., A. H. G. Rinnooy Kan, P. Brucker. 1977. Complexity of machine scheduling problems. *Annals of Discrete Mathematics* **1** 343–362.
- [14] Megow, N. 2013. Approximation in stochastic scheduling. Invited talk at the Dagstuhl Seminar 13111 “Scheduling”, Schloss Dagstuhl, Wadern, 2013.
- [15] Megow, N., M. Uetz, T. Vredeveld. 2006. Models and algorithms for stochastic online scheduling. *Mathematics of Operations Research* **31**(3) 513–525.

- [16] Megow, N., T. Vredeveld. 2014. A tight 2-approximation for preemptive stochastic scheduling. *Math. Oper. Res.* **39**(4) 1297–1310.
- [17] Möhring, R. H., F. J. Radermacher, G. Weiss. 1984. Stochastic scheduling problems I: General strategies. *ZOR - Zeitschrift für Operations Research* **28** 193–260.
- [18] Möhring, R. H., A. S. Schulz, M. Uetz. 1999. Approximation in stochastic scheduling: The power of LP-based priority policies. *Journal of the ACM* **46** 924–942.
- [19] Motwani, R., P. Raghavan. 1995. *Randomized Algorithms*. Cambridge University Press.
- [20] Queyranne, M. 1993. Structure of a simple scheduling polyhedron. *Mathematical Programming* **58** 263–285.
- [21] Ross, S. H. 1983. *Introduction to Stochastic Dynamic Programming*. Probability and Mathematical Statistics, Academic Press, New York.
- [22] Rothkopf, M. H. 1966. Scheduling with random service times. *Management Science* **12** 703–713.
- [23] Scharbrodt, M., T. Schickinger, A. Steger. 2006. A new average case analysis for completion time scheduling. *Journal of the ACM* **53** 121–146.
- [24] Schulz, A. S. 1996. Scheduling to minimize total weighted completion time: Performance guarantees of LP-based heuristics and lower bounds. W. H. Cunningham, S. T. McCormick, M. Queyranne, eds., *Integer Programming and Combinatorial Optimization, Lecture Notes in Computer Science*, vol. 1084. Springer, 301–315.
- [25] Schulz, A. S. 2008. Stochastic online scheduling revisited. B. Yang, D.-Z. Du, C. Wang, eds., *Combinatorial Optimization and Applications, Lecture Notes in Computer Science*, vol. 5165. Springer, 448–457.
- [26] Schulz, A. S., M. Skutella. 2002. Scheduling unrelated machines by randomized rounding. *SIAM Journal on Discrete Mathematics* **15** 450–469.
- [27] Schuurman, P., G. J. Woeginger. 1999. Polynomial time approximation algorithms for machine scheduling: Ten open problems. *Journal of Scheduling* **2** 203–213.
- [28] Sethuraman, J., M. S. Squillante. 1999. Optimal scheduling of multiclass parallel machines. *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms*. 963–964.
- [29] Skutella, M. 1998. Approximation and randomization in scheduling. Ph.D. thesis, Technische Universität Berlin, Germany.
- [30] Skutella, M. 2001. Convex quadratic and semidefinite programming relaxations in scheduling. *Journal of the ACM* **48** 206–242.
- [31] Skutella, M., M. Sviridenko, M. Uetz. 2014. Stochastic scheduling on unrelated machines. N. Portier E. W. Mayr, ed., *Proc. 31st Symp. Theoretical Aspects of Computer Science, Leibniz International Proceedings in Informatics*, vol. 25. Dagstuhl Publishing, Germany, 639–650.
- [32] Skutella, M., M. Uetz. 2005. Stochastic machine scheduling with precedence constraints. *SIAM Journal on Computing* **34** 788–802.
- [33] Skutella, M., G. J. Woeginger. 2000. A PTAS for minimizing the total weighted completion time on identical parallel machines. *Mathematics of Operations Research* **25** 63–75.
- [34] Smith, W. E. 1956. Various optimizers for single-stage production. *Naval Research and Logistics Quarterly* **3** 59–66.
- [35] Souza, A., A. Steger. 2006. The expected competitive ratio for weighted completion time scheduling. *Theory of Computing Systems* **39** 121–136.
- [36] Uetz, M. 2003. When greediness fails: Examples from stochastic scheduling. *Operations Research Letters* **31** 413–419.
- [37] Weiss, Gideon. 1990. Approximation results in parallel machines stochastic scheduling. *Annals of Operations Research* **26** 195–242.
- [38] Weiss, Gideon. 1992. Turnpike optimality of Smith’s rule in parallel machines stochastic scheduling. *Mathematics of Operations Research* **17** 255–270.

- [39] Wolsey, L. A. 1985. Mixed integer programming formulations for production planning and scheduling problems. Invited talk at the 12th International Symposium on Mathematical Programming, MIT, Cambridge, 1985.