

Generalizing the Kawaguchi-Kyan Bound to Stochastic Parallel Machine Scheduling

Sven Jäger

Institut für Mathematik, Technische Universität Berlin, Germany
jaeger@math.tu-berlin.de

Martin Skutella

Institut für Mathematik, Technische Universität Berlin, Germany
skutella@math.tu-berlin.de

Abstract

Minimizing the sum of weighted completion times on m identical parallel machines is one of the most important and classical scheduling problems. For the stochastic variant where processing times of jobs are random variables, Möhring, Schulz, and Uetz (1999) presented the first and still best known approximation result, achieving, for arbitrarily many machines, performance ratio $1 + \frac{1}{2}(1 + \Delta)$, where Δ is an upper bound on the squared coefficient of variation of the processing times. We prove performance ratio $1 + \frac{1}{2}(\sqrt{2} - 1)(1 + \Delta)$ for the same underlying algorithm—the Weighted Shortest Expected Processing Time (WSEPT) rule. For the special case of deterministic scheduling (i.e., $\Delta = 0$), our bound matches the tight performance ratio $\frac{1}{2}(1 + \sqrt{2})$ of this algorithm (WSPT rule), derived by Kawaguchi and Kyan in a 1986 landmark paper. We present several further improvements for WSEPT’s performance ratio, one of them relying on a carefully refined analysis of WSPT yielding, for every fixed number of machines m , WSPT’s exact performance ratio of order $\frac{1}{2}(1 + \sqrt{2}) - O(1/m^2)$.

2012 ACM Subject Classification Mathematics of computing → Combinatorial optimization, Theory of computation → Scheduling algorithms

Keywords and phrases Stochastic Scheduling, Parallel Machines, Approximation Algorithm, List Scheduling, Weighted Shortest (Expected) Processing Time Rule

Digital Object Identifier 10.4230/LIPIcs.STACS.2018.43

Related Version A full version of this paper can be found at <https://arxiv.org/abs/1801.01105>.

Funding This research was carried out in the framework of MATHEON and supported by the Einstein Foundation Berlin.

Acknowledgements We would like to thank the anonymous referees for careful reading and helpful comments.

1 Introduction

In an archetypal machine scheduling problem, n independent jobs have to be scheduled on m identical parallel machines or processors. Each job j is specified by its processing time $p_j > 0$ and by its weight $w_j > 0$. In a feasible schedule, every job j is processed for p_j time units on one of the m machines in an uninterrupted fashion, and every machine can process at most one job at a time. The completion time of job j in some schedule S is denoted by C_j^S . The goal is to compute a schedule S that minimizes the total weighted completion time



© Sven Jäger and Martin Skutella;
licensed under Creative Commons License CC-BY
35th Symposium on Theoretical Aspects of Computer Science (STACS 2018).
Editors: Rolf Niedermeier and Brigitte Vallée; Article No. 43; pp. 43:1–43:14
Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



SYMPOSIUM
ON THEORETICAL
ASPECTS
OF COMPUTER
SCIENCE


$\sum_{j=1}^n w_j C_j^S$. In the standard classification scheme of Graham, Lawler, Lenstra, and Rinnooy Kan [7], this NP-hard scheduling problem is denoted by $P||\sum w_j C_j$.

Weighted Shortest Processing Time Rule. By a well-known result of Smith [24], sequencing the jobs in order of non-increasing ratios w_j/p_j gives an optimal single-machine schedule. List scheduling in this order is known as the Weighted Shortest Processing Time (WSPT) rule and can also be applied to identical parallel machines, where it is a $\frac{1}{2}(1 + \sqrt{2})$ -approximation algorithm; see Kawaguchi and Kyan [14]. A particularly remarkable aspect of Kawaguchi and Kyan's work is that, in contrast to the vast majority of approximation results, their analysis does not rely on some kind of lower bound. Instead, they succeed in explicitly identifying a class of worst-case instances. In particular, the performance ratio $\frac{1}{2}(1 + \sqrt{2})$ is tight: For every $\varepsilon > 0$ there is a problem instance for which WSPT has approximation ratio at least $\frac{1}{2}(1 + \sqrt{2}) - \varepsilon$. The instances achieving these approximation ratios, however, have large numbers of machines when ε becomes small. Schwiegelshohn [20] gives a considerably simpler version of Kawaguchi and Kyan's analysis.

Stochastic Scheduling. Many real-world machine scheduling problems exhibit a certain degree of uncertainty about the jobs' processing times. This characteristic is captured by the theory of stochastic machine scheduling, where the processing time of job j is no longer a given number p_j but a random variable \mathbf{p}_j . As all previous work in the area, we always assume that these random variables are stochastically independent. At the beginning, only the distributions of these random variables are known. The actual processing time of a job becomes only known upon its completion. As a consequence, the solution to a stochastic scheduling problem is no longer a simple schedule, but a so-called *non-anticipative scheduling policy*. Precise definitions on stochastic scheduling policies are given by Möhring, Radermacher, and Weiss [16]. Intuitively, whenever a machine is idle at time t , a non-anticipative scheduling policy may decide to start a job of its choice based on the observed past up to time t as well as the a priori knowledge of the jobs processing time distributions and weights. It is, however, not allowed to anticipate information about the future, i.e., the actual realizations of the processing times of jobs that have not yet finished by time t .

It follows from simple examples that, in general, a non-anticipative scheduling policy cannot yield an optimal schedule for each possible realization of the processing times. We are therefore looking for a policy which minimizes the objective in expectation. For the stochastic scheduling problem considered in this paper, the goal is to find a non-anticipative scheduling policy that minimizes the expected total weighted completion time. This problem is denoted by $P|\mathbf{p}_j \sim \text{stoch}|\mathbb{E}[\sum w_j C_j]$.

Weighted Shortest Expected Processing Time Rule. The stochastic analogue of the WSPT rule is greedily scheduling the jobs in order of non-increasing ratios $w_j/\mathbb{E}[\mathbf{p}_j]$. Whenever a machine is idle, the Weighted Shortest Expected Processing Time (WSEPT) rule immediately starts the next job in this order. For a single machine this is again optimal; see Rothkopf [18]. For identical parallel machines, Cheung, Fischer, Matuschke, and Megow [3] and Im, Moseley, and Pruhs [10] independently show that WSEPT does not even achieve constant performance ratio. More precisely, for every $R > 0$ there is a problem instance for which WSEPT's expected total weighted completion time is at least R times the expected objective value of an optimal non-anticipative scheduling policy. In the special case of exponentially distributed processing times, Jagtenberg, Schwiegelshohn, and Uetz [12] show a lower bound of 1.243 on WSEPT's performance. On the positive side, WSEPT is an optimal



■ **Figure 1** Bounds on WSEPT's performance ratio.


policy for the special case of unit weight jobs with stochastically ordered processing times, $P|\mathbf{p}_j \sim \text{stoch}(\preceq_{\text{st}})|E[\sum C_j]$; see Weber, Varaiya, and Walrand [25]. Moreover, Weiss [26, 27] proves asymptotic optimality of WSEPT for bounded second moments of the residual processing time distributions. Möhring, Schulz, and Uetz [17] show that WSEPT achieves performance ratio $1 + \frac{1}{2}(1 + \Delta)(1 - \frac{1}{m})$, where Δ is an upper bound on the squared coefficient of variation of the processing times.

Further Approximation Results from the Literature. While there is a PTAS for the deterministic problem $P||\sum w_j C_j$ [23], no constant-factor approximation algorithm is known for the stochastic problem $P|\mathbf{p}_j \sim \text{stoch}|E[\sum w_j C_j]$. WSEPT's performance ratio $1 + \frac{1}{2}(1 + \Delta)$ (for arbitrarily many machines) proven by Möhring et al. [17] is the best hitherto known performance ratio. The only known approximation ratio not depending on the jobs' squared coefficient of variation Δ is due to Im et al. [10], who, for the special case of unit job weights $P|\mathbf{p}_j \sim \text{stoch}|E[\sum C_j]$, present an $O(\log^2 n + m \log n)$ -approximation algorithm.

The performance ratio $1 + \frac{1}{2}(1 + \Delta)$ has been carried over to different generalizations of $P|\mathbf{p}_j \sim \text{stoch}|E[\sum w_j C_j]$. Megow, Uetz, and Vredeveld [15] show that it also applies if jobs arrive online in a list and must immediately and irrevocably be assigned to machines, on which they can be sequenced optimally. An approximation algorithm with this performance ratio for the problem on unrelated parallel machines is designed by Skutella, Sviridenko, and Uetz [22]. If these two features are combined, i.e., in the online list-model with unrelated machines, Gupta, Moseley, Uetz, and Xie [8] develop a $(8 + 4\Delta)$ -approximation algorithm.

The performance ratios are usually larger if jobs are released over time: In the offline setting with identical machines the best known approximation algorithm has performance ratio $2 + \Delta$; see Schulz [19]. This performance ratio is also achieved for unrelated machines [22] and by a randomized online algorithm [19]. In the online setting there exist furthermore a deterministic $(\max\{2.618, 2.309 + 1.309\Delta\})$ -approximation on identical machines [19] and a deterministic $(144 + 72\Delta)$ -approximation on unrelated machines [8].

Our Contribution and Outline. We present the first progress on the approximability of the basic stochastic scheduling problem on identical parallel machines with expected total weighted completion time objective $P|\mathbf{p}_j \sim \text{stoch}|E[\sum w_j C_j]$ since the seminal work of Möhring et al. [17]; see Figure 1. We prove that WSEPT achieves performance ratio



■ **Figure 2** Graph of the function $m \mapsto 1 + \frac{1}{2}(\sqrt{(2m - k_m)k_m} - k_m)/m$, which for $m \in \mathbb{N}$ gives the worst-case approximation ratio of WSPT for $P||\sum w_j C_j$ with m machines (dots), compared to the machine-independent Kawaguchi-Kyan bound.

$$1 + \frac{1}{2} \min \left\{ \frac{\sqrt{(2m - k_m)k_m} - k_m}{m}, \frac{1}{1 + \min\{2, \sqrt{2 + 2\Delta}\}} \right\} (1 + \Delta), \quad (1)$$

where $k_m := \lfloor (1 - \frac{1}{2}\sqrt{2})m \rfloor$ is the nearest integer to $(1 - \frac{1}{2}\sqrt{2})m$. Notice that, for every number of machines m , the performance ratio given by the first term of the minimum in (1) is bounded from above by $1 + \frac{1}{2}(\sqrt{2} - 1)(1 + \Delta)$, and for $m \rightarrow \infty$ it converges to this bound. As $(1 + \min\{2, \sqrt{2 + 2\Delta}\})^{-1} \leq \sqrt{2} - 1$ for all $\Delta > 0$, when considering an arbitrary number of machines, the second term in the minimum dominates the first term. In the following, we list several points that emphasize the significance of the new performance ratio (1).

- For the special case of deterministic scheduling (i.e., $\Delta = 0$), the machine-independent performance ratio in (1) matches the Kawaguchi-Kyan bound $\frac{1}{2}(1 + \sqrt{2})$, which is known to be tight [14]. In particular, we dissolve the somewhat annoying discontinuity of the best previously known bounds [14, 17] at $\Delta = 0$; see Figure 1.
- Again for deterministic jobs, our machine-dependent bound $1 + \frac{1}{2}(\sqrt{(2m - k_m)k_m} - k_m)/m$ is tight and slightly improves the 30 years old Kawaguchi-Kyan bound for every fixed number of machines m ; see Figure 2.
- For exponentially distributed processing times ($\Delta = 1$), our results imply that WSEPT achieves performance ratio $4/3$. This solves an open problem by Jagtenberg et al. [12], who give a lower bound of 1.243 on WSEPT's performance and ask for an improvement of the previously best known upper bound of $2 - 1/m$ due to Möhring et al. [17].
- WSEPT's performance bound due to Möhring et al. [17] also holds for the MinIncrease policy, introduced by Megow et al. [15], which is a fixed-assignment policy, i.e. it determines for each job beforehand on which machine it is processed. Our stronger bound, together with a lower bound in [22], shows that WSEPT beats every fixed-assignment policy.

The improved performance ratio in (1) is derived as follows. In Section 2 we present one of the key results of this paper (see Theorem 1 below): If WSPT has performance ratio $1 + \beta$ for some β , then WSEPT achieves performance ratio $1 + \beta(1 + \Delta)$ for the stochastic scheduling problem. For the Kawaguchi-Kyan bound $1 + \beta = \frac{1}{2}(1 + \sqrt{2})$, this yields performance ratio $1 + \frac{1}{2}(\sqrt{2} - 1)(1 + \Delta)$. It is also interesting to notice that the performance ratio of Möhring et al. [17] follows from this theorem by plugging in $1 + \beta = 3/2 - 1/(2m)$, which is WSPT's performance ratio obtained from the bound of Eastman, Even, and Isaacs [4]; see Kawaguchi and Kyan [14]. We generalize Theorem 1 to performance ratios w.r.t. the weighted sum of α -points as objective function, where the α -point of a job j is the point in time when it has been processed for exactly αp_j time units.

The theorems derived in Section 2 provide tools to carry over bounds for the WSPT rule to the WSEPT rule. The concrete performance ratio for the WSEPT rule obtained this way thus depends on good bounds for the WSPT rule. In Section 3 we derive performance ratios for WSPT w.r.t. the weighted sum of α -points objective. For $\alpha = \frac{1}{2}$ this performance ratio follows easily from a result by Avidor, Azar, and Sgall [1]. As a consequence we obtain performance ratio $1 + \frac{1}{6} \max\{2, 1 + \Delta\}$ for WSEPT. By optimizing the choice of α , we finally obtain the performance ratio $1 + \frac{1}{2}(1 + \min\{2, \sqrt{2 + 2\Delta}\})^{-1}(1 + \Delta)$. The various bounds derived in Sections 2 and 3 are illustrated in Figure 1. Finally, in Section 4 the analysis of Schwegelshohn [20] for the WSPT rule is refined for every fixed number of machines m , entailing the machine-dependent bound for the WSEPT rule in (1).

Due to space constraints, some proofs are omitted in this extended abstract. They can be found in the full version of this paper [11].

2 Performance ratio of the WSEPT rule

Let $\Delta \geq \text{Var}[p_j]/\mathbb{E}[p_j]^2$ for all $j \in \{1, \dots, n\}$. In Theorems 1 and 3 we demonstrate how performance ratios for the WSPT rule for deterministic scheduling can be carried over to stochastic scheduling. Theorem 1 starts out from a performance ratio for WSPT with respect to the usual objective function: the weighted sum of completion times. In Theorem 3 this is generalized insofar as a performance ratio for WSPT with the weighted sum of α -points as objective function is taken as a basis. Only Theorem 1 is proven in this extended abstract.

► **Theorem 1.** *If the WSPT rule on m machines has performance ratio $1 + \beta_m$ for the problem $P \parallel \sum w_j C_j$, then the WSEPT rule achieves performance ratio $1 + \beta_m(1 + \Delta)$ for $P \parallel p_j \sim \text{stoch} \mid \mathbb{E}[\sum w_j C_j]$ on m machines.*

The reason why the bound for the WSPT rule does not directly carry over to the WSEPT rule is that under a specific realization of the processing times the schedule obtained by the WSEPT policy may differ from the WSPT schedule for this realization. Still, under every realization the WSEPT schedule is a list schedule. Hence, usually a bound that is valid for every list schedule is used: The objective value of a list schedule on m machines is at most $1/m$ times the objective value of the list schedule on a single machine plus $(m - 1)/m$ times the weighted sum of processing times. This bound, holding because a list scheduling policy assigns each job to the currently least loaded machine, is applied realizationwise to obtain a corresponding bound on the expected values in stochastic scheduling (cf. [17, Lemma 4.1]), which is then compared to an LP-based lower bound on the expected total weighted completion time under an optimal scheduling policy.

In order to benefit from the precise bounds known for the WSPT rule nevertheless, we regard the following auxiliary stochastic scheduling problem: For each job, instead of its weight w_j , we are given a weight factor ρ_j . The actual weight of a job is ρ_j times its actual processing time, i.e., if a job takes longer, it also becomes more important. The goal is again to minimize the total weighted completion time. For the thus defined stochastic scheduling problem list scheduling in order of the ρ_j has the nice property that it creates a WSPT schedule in every realization. So, any performance ratio of the WSPT rule directly carries over to this list scheduling policy for the auxiliary scheduling problem. In the following proof of Theorem 1 we first compare the expected total weighted completion time of a WSEPT schedule for the original problem to the expected objective value of the schedule obtained by list scheduling in order of ρ_j for the auxiliary problem, then apply the performance ratio of the WSPT rule, and finally compare the expected total weighted completion time of an

optimal schedule for the auxiliary problem to the expected objective value of the schedule obtained by an optimal policy for the original problem. The transitions between the two problems lead to the additional factor $1 + \Delta$ in the performance ratio.

Proof. Consider an instance of $P|\mathbf{p}_j \sim \text{stoch}|E[\sum w_j C_j]$ consisting of n jobs and m machines, and let $\beta := \beta_m$ and $\rho_j := w_j/E[\mathbf{p}_j]$ for $j \in \{1, \dots, n\}$. For every realization $\vec{p} = (p_1, \dots, p_n)$ of the processing times we consider the instance $I(\vec{p})$ of $P|\sum w_j C_j$ which consists of n jobs with processing times p_1, \dots, p_n and weights $\rho_1 p_1, \dots, \rho_n p_n$, so that the jobs in this instance have Smith ratios ρ_1, \dots, ρ_n under all possible realizations. Therefore, for every realization \vec{p} the schedule obtained by the WSEPT policy is a WSPT schedule for $I(\vec{p})$. Let $C_j^{\text{WSEPT}}(\vec{p})$ denote the completion time of job j in the schedule obtained by the WSEPT policy in the realization \vec{p} , let $C_j^*(I(\vec{p}))$ denote its completion time in an optimal schedule for $I(\vec{p})$, and let $C_j^{\Pi^*}(\vec{p})$ denote j 's completion time in the schedule constructed by an optimal stochastic scheduling policy under the realization \vec{p} . For every realization \vec{p} of the processing times, since the WSEPT schedule obeys the WSPT rule for $I(\vec{p})$, its objective value is bounded by

$$\sum_{j=1}^n (\rho_j p_j) C_j^{\text{WSEPT}}(\vec{p}) \leq (1 + \beta) \cdot \sum_{j=1}^n (\rho_j p_j) C_j^*(I(\vec{p})).$$

As the schedule obtained by an optimal stochastic scheduling policy is feasible for $I(\vec{p})$,

$$\sum_{j=1}^n (\rho_j p_j) C_j^*(I(\vec{p})) \leq \sum_{j=1}^n (\rho_j p_j) C_j^{\Pi^*}(\vec{p}).$$

By putting these two inequalities together and taking expectations, we get the inequality

$$\mathbb{E} \left[\sum_{j=1}^n \rho_j \mathbf{p}_j C_j^{\text{WSEPT}} \right] \leq (1 + \beta) \cdot \mathbb{E} \left[\sum_{j=1}^n \rho_j \mathbf{p}_j C_j^{\Pi^*} \right],$$

where $C_j^{\text{WSEPT}} = C_j^{\text{WSEPT}}((\mathbf{p}_1, \dots, \mathbf{p}_n))$ and $C_j^{\Pi^*} = C_j^{\Pi^*}((\mathbf{p}_1, \dots, \mathbf{p}_n))$. Using the latter inequality, we can bound the expected total weighted completion time of the WSEPT rule:

$$\begin{aligned} & \mathbb{E} \left[\sum_{j=1}^n w_j C_j^{\text{WSEPT}} \right] = \sum_{j=1}^n \rho_j \mathbb{E}[\mathbf{p}_j] \mathbb{E}[C_j^{\text{WSEPT}}] \\ & \stackrel{(*)}{=} \sum_{j=1}^n \rho_j \mathbb{E}[\mathbf{p}_j C_j^{\text{WSEPT}}] - \sum_{j=1}^n \rho_j \text{Var}[\mathbf{p}_j] = \mathbb{E} \left[\sum_{j=1}^n \rho_j \mathbf{p}_j C_j^{\text{WSEPT}} \right] - \sum_{j=1}^n \rho_j \text{Var}[\mathbf{p}_j] \\ & \leq (1 + \beta) \mathbb{E} \left[\sum_{j=1}^n \rho_j \mathbf{p}_j C_j^{\Pi^*} \right] - \sum_{j=1}^n \rho_j \text{Var}[\mathbf{p}_j] = (1 + \beta) \sum_{j=1}^n \rho_j \mathbb{E}[\mathbf{p}_j C_j^{\Pi^*}] - \sum_{j=1}^n \rho_j \text{Var}[\mathbf{p}_j] \\ & \stackrel{(*)}{=} (1 + \beta) \cdot \left(\sum_{j=1}^n \rho_j \mathbb{E}[\mathbf{p}_j] \mathbb{E}[C_j^{\Pi^*}] + \sum_{j=1}^n \rho_j \text{Var}[\mathbf{p}_j] \right) - \sum_{j=1}^n \rho_j \text{Var}[\mathbf{p}_j] \\ & = (1 + \beta) \cdot \sum_{j=1}^n w_j \mathbb{E}[C_j^{\Pi^*}] + \beta \sum_{j=1}^n \rho_j \text{Var}[\mathbf{p}_j] \leq (1 + \beta) \cdot \sum_{j=1}^n w_j \mathbb{E}[C_j^{\Pi^*}] + \Delta \beta \sum_{j=1}^n w_j \mathbb{E}[\mathbf{p}_j] \\ & \leq (1 + \beta(1 + \Delta)) \cdot \sum_{j=1}^n w_j \mathbb{E}[C_j^{\Pi^*}] = (1 + \beta(1 + \Delta)) \cdot \mathbb{E} \left[\sum_{j=1}^n w_j C_j^{\Pi^*} \right]. \end{aligned}$$

The equalities marked with (*) hold because for any stochastic scheduling policy Π and all j

$$E[\mathbf{p}_j \mathbf{C}_j^\Pi] = E[\mathbf{p}_j \mathbf{S}_j^\Pi] + E[\mathbf{p}_j^2] = E[\mathbf{p}_j]E[\mathbf{S}_j^\Pi] + E[\mathbf{p}_j]^2 + \text{Var}[\mathbf{p}_j] = E[\mathbf{p}_j]E[\mathbf{C}_j^\Pi] + \text{Var}[\mathbf{p}_j],$$

where \mathbf{S}_j^Π denotes the starting time of job j under policy Π . The independence of \mathbf{p}_j and \mathbf{S}_j^Π follows from the independence of the processing times and the non-anticipativity of policy Π , and the last inequality uses the fact that $E[\mathbf{p}_j] \leq E[\mathbf{C}_j^{\Pi^*}]$ for every job j . ◀

By plugging in the Kawaguchi-Kyan bound, we immediately get the following performance ratio (see Figure 1).

► **Corollary 2.** *The WSEPT rule has performance ratio $1 + \frac{1}{2}(\sqrt{2} - 1) \cdot (1 + \Delta)$ for the problem $P|\mathbf{p}_j \sim \text{stoch}|E[\sum w_j \mathbf{C}_j]$.*

For $\alpha \in (0, 1]$ the α -point $C_j^S(\alpha)$ of a job j is the (first) point in time at which it has been processed for αp_j time units. Introduced by Hall, Shmoys, and Wein [9] in order to convert a preemptive schedule into a non-preemptive one, the concept of α -points is often used in the *design* of algorithms (see e.g. [5, 2, 6, 21]). In contrast, we use them in the definition of an alternative objective function in order to improve the *analysis* of the WSEPT rule.

We consider as objective function the weighted sum of α -points $\sum_{j=1}^n w_j C_j^S(\alpha)$ for $\alpha \in (0, 1]$. This differs only by the constant $(1 - \alpha) \sum_{j=1}^n w_j p_j$ from the weighted sum of completion times. So as for optimal solutions the objective functions are equivalent. The same applies to the stochastic variant, where the two objectives differ by $(1 - \alpha) \sum_{j=1}^n w_j E[\mathbf{p}_j]$. We now generalize Theorem 1 to the (expected) weighted sum of α -points.

► **Theorem 3.** *If the WSPT rule has performance ratio $1 + \beta$ for the deterministic problem $P|\sum w_j C_j(\alpha)$, then the WSEPT rule has performance ratio $1 + \beta(1 + \Delta)$ for the problem $P|\mathbf{p}_j \sim \text{stoch}|E[\sum w_j \mathbf{C}_j(\alpha)]$ and $1 + \beta \cdot \max\{1, \alpha(1 + \Delta)\}$ for $P|\mathbf{p}_j \sim \text{stoch}|E[\sum w_j \mathbf{C}_j]$.*

The proof relies on the same idea as the proof of Theorem 1, namely to apply the bound for $P|\sum w_j C_j(\alpha)$ realizationwise to the auxiliary stochastic problem described above. Theorem 1 follows from Theorem 3 by plugging in $\alpha = 1$.

3 Performance ratios for WSPT with weighted sum of α -points objective

In this section we derive performance ratios for $P|\sum w_j C_j(\alpha)$. The two classical performance guarantees for $P|\sum w_j C_j$ by Eastman, Even, and Isaacs [4] and by Kawaguchi and Kyan [14] can both be generalized to this problem. While the Eastman-Even-Isaacs bound can be established for every $\alpha \in (0, 1]$, the Kawaguchi-Kyan bound carries over only for $\alpha \in [\frac{1}{2}, 1]$. In return, the generalized Kawaguchi-Kyan bound is better for these α .

For a problem instance I denote by $\mathcal{N}(I)$ its job set, by $C_j^{\text{WSPT}}(\alpha)(I)$ the α -point of job j in the WSPT schedule for I , and by $C_j^*(\alpha)(I)$ the α -point of job j in some fixed (‘the’) optimal schedule for I . Hence $C_j^{\text{WSPT}}(1)(I) = C_j^{\text{WSPT}}(I)$ is the completion time of j in the WSPT schedule, and analogously for the optimal schedule. Furthermore, let $M_i^{\text{WSPT}}(I)$ and $M_i^*(I)$ denote the load of the i -th machine and $M_{\min}^{\text{WSPT}}(I)$ and $M_{\min}^*(I)$ denote the load of the least loaded machine, in the WSPT schedule and the optimal schedule for I , respectively. Moreover, let $\text{WSPT}_\alpha(I)$ and $\text{OPT}_\alpha(I)$ denote the weighted sum of α -points of the schedule obtained by the WSPT rule and of the optimal schedule, respectively. Finally, denote by $\lambda_\alpha(I) := \text{WSPT}_\alpha(I)/\text{OPT}_\alpha(I)$ the approximation ratio of the WSPT rule for the instance I . We assume that if multiple jobs have the same ratio w_j/p_j , the WSPT rule processes them according to an arbitrary job order given as part of the input.

It is a well-known fact (see e.g. [20]) that for the weighted sum of completion times objective the worst case for the WSPT rule occurs if all jobs have the same Smith ratio w_j/p_j . This generalizes to the weighted sum of α -points objective.

► **Lemma 4.** *For every $\alpha \in [0, 1]$ and every instance I of $P||\sum w_j C_j(\alpha)$ there is an instance I' of $P||\sum p_j C_j(\alpha)$ with the same number of machines and $\lambda_\alpha(I') \geq \lambda_\alpha(I)$.*

The proof proceeds in the same way as the proof of Schwiegelshohn [20]. For unit Smith ratio instances the WSPT rule is nothing but list scheduling according to an arbitrary given order. Restricting to them has the benefit that the objective value of a schedule S can be computed easily from its machine loads, namely

$$\sum_{j=1}^n p_j C_j^S(\tfrac{1}{2}) = \frac{1}{2} \sum_{i=1}^m (M_i^S)^2. \quad (2)$$

This classical observation can for example be found in the paper of Eastman et al. [4].

For the sum of the squares of the machine loads as objective function Avidor, Azar, and Sgall [1] showed that WSPT has performance ratio $4/3$. So this also holds for the weighted sum of $\frac{1}{2}$ -points. By plugging it in into Theorem 3, we get the following corollary.

► **Corollary 5.** *The WSEPT rule has performance ratio $1 + \frac{1}{6} \max\{2, 1 + \Delta\}$ for the scheduling problem $P|p_j \sim \text{stoch}|E[\sum w_j C_j]$.*

Now we generalize the bound of Eastman, Even, and Isaacs [4].

► **Theorem 6** (Generalized Eastman-Even-Isaacs bound). *For every $\alpha \in (0, 1]$ the WSPT rule has performance ratio*

$$1 + \frac{m-1}{2\alpha m} \leq 1 + \frac{1}{2\alpha}$$

for the problem $P||\sum w_j C_j(\alpha)$.

► **Remark.** The generalized Eastman-Even-Isaacs bound does not lead to better performance ratios for the WSEPT rule for $P|p_j \sim \text{stoch}|E[\sum w_j C_j]$ than the bound of Möhring et al. [17], as plugging in $\beta = \frac{m-1}{2\alpha m}$ into Theorem 3 leads to a performance ratio of

$$1 + \frac{m-1}{2\alpha m} \cdot \max\{1, \alpha(1 + \Delta)\} \geq 1 + \frac{1}{2}(1 + \Delta) \left(1 - \frac{1}{m}\right).$$

So far, by choosing $\alpha = 1$ and $\alpha = \frac{1}{2}$ we have derived the two performance ratios for the WSEPT rule labeled by [Cor. 2] and [Cor. 5] in Figure 1. The proofs of Schwiegelshohn [20] and of Avidor et al. [1] of the underlying bounds for WSPT are quite similar. Both consist of a sequence of steps that reduce the set of instances to be examined. In every such reduction step it is shown that for any instance I of the currently considered set there is an instance I' in a smaller set for which the approximation ratio of WSPT is not better. This can be generalized to arbitrary $\alpha \in [\frac{1}{2}, 1]$. The resulting performance ratios for WSPT lead by means of Theorem 3 to a family of different performance ratios for the WSEPT rule. Note that the performance ratio of WSEPT following from the result of Avidor et al. for $\alpha = \frac{1}{2}$ has better behavior for large values of Δ , while the performance ratio following from Kawaguchi and Kyan's result for $\alpha = 1$ is better for small Δ . This behavior generalizes to $\alpha \in [\frac{1}{2}, 1]$: the smaller the underlying α , the better the ratio for large Δ but the worse the ratio for small Δ . Finally, we take for every $\Delta > 0$ the minimum of all the derived bounds.

► **Theorem 7** (Generalized Kawaguchi-Kyan bound). *For every $\alpha \in [\frac{1}{2}, 1]$ the WSPT rule has performance ratio*

$$1 + \frac{1}{2\alpha + \sqrt{8\alpha}}$$

for $P \parallel \sum w_j C_j(\alpha)$, and this bound is tight.

Combining this bound with Theorem 3 yields for every $\alpha \in [\frac{1}{2}, 1]$ the performance ratio $1 + \frac{1}{2} \max\{1/(\alpha + \sqrt{2\alpha}), (1 + \Delta)/(1 + \sqrt{2/\alpha})\}$ of WSEPT for $P \parallel \sum w_j C_j$. This is minimized at $\alpha := 1/\min\{2, 1 + \Delta\}$, yielding the following performance ratio (see Figure 1).

► **Corollary 8.** *For $P \parallel \sum w_j C_j$ the WSEPT rule has performance ratio*

$$1 + \frac{1}{2} \cdot \frac{1}{1 + \min\{2, \sqrt{2(1 + \Delta)}\}} \cdot (1 + \Delta).$$

Proof sketch of Theorem 7

The proof of Theorem 7 is analogous to the proof of Schwiegelshohn [20], consisting of a sequence of reduction lemmas. Let $\alpha \in [\frac{1}{2}, 1]$, assume that $p_1 \geq \dots \geq p_n$, and let $\ell := \max\{j \in \{1, \dots, m\} \mid p_j \geq \frac{1}{m-j+1} \sum_{j'=j}^n p_{j'}\}$. Then we call the ℓ jobs with largest processing times *long* jobs and denote the set of long jobs by \mathcal{L} .

► **Lemma 9.** *For every instance I of $P \parallel \sum p_j C_j(\alpha)$ and every $\varepsilon > 0$ there is an instance $I' = I'(\varepsilon)$ of $P \parallel \sum p_j C_j(\alpha)$ with the same number of machines such that $\lambda_\alpha(I') \geq \lambda_\alpha(I)$ and*

1. $M_{\min}^{\text{WSPT}}(I') = 1$,
2. every job j with $S_j^{\text{WSPT}}(I') < M_{\min}^{\text{WSPT}}(I')$ fulfills $C_j^{\text{WSPT}}(I') \leq M_{\min}^{\text{WSPT}}(I')$ and $p'_j < \varepsilon$,
3. in the optimal schedule for I' every machine is used only by a single long job or has load $M_{\min}^*(I')$.


Like in Schwiegelshohn's paper, the lemma is proven by scaling the instance and splitting all jobs with $S_j^{\text{WSPT}} < M_{\min}^{\text{WSPT}}$ until they satisfy the conditions. Note that the restriction to $\alpha \geq \frac{1}{2}$ is needed for this lemma because for smaller α splitting jobs increases the objective value and can thence reduce the performance ratio.

From now on, we focus on instances I that fulfill the requirements of Lemma 9 for some $0 < \varepsilon < M_{\min}^{\text{WSPT}}(I)$. For a subset $\mathcal{J} \subseteq \mathcal{N}$ of jobs we write $p(\mathcal{J}) := \sum_{j \in \mathcal{J}} p_j$. We call the jobs in $\mathcal{S} := \{j \in \{1, \dots, n\} \mid S_j^{\text{WSPT}}(I) < M_{\min}^{\text{WSPT}}\}$ *short* jobs. This set is disjoint from \mathcal{L} because all jobs in \mathcal{S} have processing time $p_j < \varepsilon$, and all jobs in \mathcal{L} have processing time $p_j \geq p_\ell \geq \frac{1}{m-\ell+1} \sum_{j'=\ell}^n p_{j'} \geq M_{\min}^{\text{WSPT}} > \varepsilon$. Finally, we call the jobs in $\mathcal{M} := \mathcal{N}(I) \setminus (\mathcal{S} \cup \mathcal{L})$ *medium* jobs. For an instance I of the type of Lemma 9, in the optimal schedule every machine that does not process a long job has load $M_{\min}^*(I) = p(\mathcal{M} \cup \mathcal{S})/(m - |\mathcal{L}|)$. We may assume that every machine processes at most one non-short job (see Figure 3).

► **Lemma 10.** *For every instance I of $P \parallel \sum p_j C_j(\alpha)$ satisfying the conditions of Lemma 9 there is an instance I' with $\lambda_\alpha(I') \geq \lambda_\alpha(I)$ that still satisfies the conditions of Lemma 9 and has the additional property that the processing times of all non-short jobs are equal.*

The proof is an adapted version of the proof of Corollary 5 in the paper of Schwiegelshohn.

Since by Lemma 9 reducing ε can only increase the approximation ratio, the worst-case approximation ratio is approached in the limit $\varepsilon \rightarrow 0$, which we will subsequently further investigate. In the limit the sum of the squared processing times of the short jobs is negligible, wherefore the limits for $\varepsilon \rightarrow 0$ of the objective values of the WSPT schedule and the optimal



■ **Figure 3** Optimal schedule and WSPT schedule for instance satisfying the conditions of Lemma 9.

schedule for an instance $I(\varepsilon)$ of the type of Lemma 10 only depend on two variables: the ratio s between the numbers of non-short jobs and machines and the duration x of the non-short jobs. The limit of the objective value of the WSPT schedule is given by

$$\lim_{\varepsilon \rightarrow 0} \text{WSPT}_\alpha(I(\varepsilon)) = \frac{m}{2} + smx(1 + \alpha x).$$

For the optimal schedule the formula depends on whether the non-short jobs are medium or long. In the first case it is given by

$$\lim_{\varepsilon \rightarrow 0} \text{OPT}_\alpha(I(\varepsilon)) = \frac{m}{2}(sx + 1)^2 + \left(\alpha - \frac{1}{2}\right)smx^2.$$

and in the second case by

$$\lim_{\varepsilon \rightarrow 0} \text{OPT}_\alpha(I(\varepsilon)) = \alpha smx^2 + \frac{m}{2(1-s)}.$$

So we have to determine the maximum of the function

$$\lambda_M(s, x) := \frac{\frac{m}{2} + smx(1 + \alpha x)}{\frac{m}{2}(sx + 1)^2 + \left(\alpha - \frac{1}{2}\right)smx^2} = \frac{2sx(\alpha x + 1) + 1}{s^2x^2 + sx((2\alpha - 1)x + 2) + 1}$$

on $\{(s, x) \mid 0 \leq s < 1, 0 \leq x \leq 1/(1-s)\}$ and the maximum of

$$\lambda_L(s, x) := \frac{\frac{m}{2} + smx(1 + \alpha x)}{\alpha smx^2 + \frac{m}{2(1-s)}} = \frac{(1-s)(2sx(\alpha x + 1) + 1)}{2\alpha s(1-s)x^2 + 1}$$


on the region $\{(s, x) \mid 0 \leq s < 1, 1/(1-s) \leq x\}$.

The partial derivative $\frac{\partial}{\partial x} \lambda_M$ is positive on the feasible region, so for every fixed s the maximum of $\lambda_M(s, \cdot)$ is attained at $x = \frac{1}{1-s}$, corresponding to the case that the non-short jobs are long. This case is also captured by the function λ_L .

For $x \rightarrow \infty$ the function λ_L converges to one. Hence, for every s the maximum of $\lambda_L(s, \cdot)$ must be attained at a finite point x . The partial derivative $\frac{\partial}{\partial x} \lambda_L$ has only one positive root, namely $x_s := (\alpha s + \sqrt{(2(1-s) + \alpha s)\alpha s}) / (2\alpha s(1-s)) > 1/(1-s)$. By plugging this in, we obtain $\lambda_L(s, x_s) = 1 + \frac{1}{2}(\sqrt{(2(1-s) + \alpha s)\alpha s} / \alpha - s)$. The only root of the derivative of the function $s \mapsto \lambda_L(s, x_s)$ that is less than 1 is $s := 1/(2 + \sqrt{2\alpha})$. Plugging this in yields the worst-case performance ratio

$$1 + \frac{1}{2\alpha + \sqrt{8\alpha}}.$$

Like the proofs of Kawaguchi and Kyan [14] and of Avidor et al. [1], this proof shows how the worst-case instances look like: They consist of short jobs of total length m and $1/(2 + \sqrt{2\alpha})m$ long jobs of length $1 + \sqrt{2/\alpha}$. For $\alpha \in \{1/2, 1\}$ we recover the worst case instances of Avidor et al. and of Kawaguchi and Kyan.



■ **Figure 4** Optimal schedule and WSPT schedule for instance after the transformation of Lemma 12.

4 Performance ratio of the WSPT rule for a fixed number of machines

In this section we analyse the WSPT rule for the problem $P||\sum w_j C_j$ with a fixed number m of machines. The problem instances of Kawaguchi and Kyan [14] whose approximation ratios converge to $(1 + \sqrt{2})/2$ consist of a set of infinitesimally short jobs with total processing time m , and a set of k jobs of length $1 + \sqrt{2}$, where $k/m \rightarrow 1 - \sqrt{2}/2$. Since $1 - \sqrt{2}/2$ is irrational, the worst case ratio can only be approached if the number of machines goes to infinity. Rounding these instances for a fixed m by choosing k as the nearest integer to $(1 - \frac{\sqrt{2}}{2})m$ (in the following denoted by $\lfloor (1 - \frac{\sqrt{2}}{2})m \rfloor$) yields at least a lower bound on the worst-case approximation ratio for $P||\sum w_j C_j$. As we will see, the worst-case instances for any fixed m look almost like that except that the length of the long jobs depends on m .

► **Theorem 11.** *For $P||\sum w_j C_j$ the WSPT rule has performance ratio*

$$1 + \frac{1}{2} \frac{\sqrt{(2m - k_m)k_m} - k_m}{m}, \quad \text{where } k_m := \left\lfloor \left(1 - \frac{\sqrt{2}}{2}\right)m \right\rfloor.$$


Moreover, this bound is tight for every fixed $m \in \mathbb{N}$.

In the remainder we summarize the proof of this theorem. Lemmas 4 and 9 hold in particular for the weighted sum of completion times. Since the described transformations do not change the number of machines, also for a fixed number m of machines the worst case occurs in an instance of the form described in Lemma 9. However, we cannot apply Lemma 10 when m is fixed because the transformation in this lemma possibly changes the number of machines. As this is not allowed in our setting, we have to find different reductions. We first reduce to instances with at most one medium job and then reduce further to instances where all long jobs have equal length. Similar reductions are also carried out by Kalaitzis, Svensson, and Tarnawski [13].

► **Lemma 12.** *For every instance I of $P||\sum p_j C_j$ satisfying the conditions of Lemma 9 there is an instance I' with the same number of machines and $\lambda(I') \geq \lambda(I)$ that still satisfies the conditions of Lemma 9 and has the additional property that there is at most one medium job.*

For the instance I shown in Figure 3 the optimal and the WSPT schedule of the reduced instance I' are shown in Figure 4.

► **Lemma 13.** *For every instance I of $P||\sum p_j C_j$ satisfying the conditions of Lemma 12 there is an instance I' with the same number of machines and $\lambda(I') \geq \lambda(I)$ that still satisfies the conditions of Lemma 12 and additionally all long jobs have equal processing time.*



■ **Figure 5** Optimal schedule and WSPT schedule for instance after the transformation of Lemma 13.

The reduction used in the proof of this lemma is illustrated in Figure 5.

As in Section 3 we will analyse the limit for $\varepsilon \rightarrow 0$. The limits of the objective values of the WSPT schedule and the optimal schedule for an instance $I(\varepsilon)$ of the type of Lemma 13 depend only on three variables: two real variables, namely the length x of the long jobs and the length y of the medium job ($y = 0$ if no medium job exists), and one integer variable: the number k of long jobs. They are given by

$$\lim_{\varepsilon \rightarrow 0} \text{WSPT}(I(\varepsilon)) = \frac{m}{2} + kx(1+x) + y(1+y),$$

$$\lim_{\varepsilon \rightarrow 0} \text{OPT}(I(\varepsilon)) = k \cdot x^2 + \frac{(m+y)^2}{2(m-k)} + \frac{y^2}{2}.$$

In Figure 6 of the full version [11] these formulas are illustrated via two-dimensional Gantt charts for the three different types of single-machine schedules used by the WSPT schedule and the optimal schedule, respectively. In order to describe a valid scheduling instance of the prescribed type, the values x , y , and k must lie in the domains

$$k \in \{0, \dots, m-1\}, \quad y \begin{cases} \in \left[0, \frac{m}{m-k-1}\right] & \text{if } k < m-1, \\ = 0 & \text{if } k = m-1, \end{cases} \quad x \in \left[\frac{y+m}{m-k}, \infty\right).$$

► **Lemma 14.** *The maximum of the ratio*

$$\lambda_m(x, y, k) := \frac{\frac{m}{2} + kx(1+x) + y(1+y)}{k \cdot x^2 + \frac{(m+y)^2}{2(m-k)} + \frac{y^2}{2}} = \frac{(m-k)(2kx^2 + 2kx + 2y^2 + 2y + m)}{(m-k)(2kx^2 + y^2) + (y+m)^2}$$

on the feasible domains is $1 + \frac{1}{2}(\sqrt{(2m-k_m)k_m} - k_m)/m$, and it is attained at

$$k_m := \left\lfloor \left(1 - \frac{1}{2}\sqrt{2}\right)m \right\rfloor, \quad y_m := 0, \quad x_m := \frac{m}{\sqrt{(2m-k_m)k_m} - k_m}.$$

The calculations leading to these values are similar to those in Section 3. This concludes the proof of Theorem 11.

In Figure 2 the function $m \mapsto \lambda_m(x_m, 0, k_m)$, whose values at integral m are exactly the worst case approximation ratios for instances with m machines, is depicted. The jumps and kinks occur when the number k_m of long jobs in the worst-case instance changes. By taking the limit for $m \rightarrow \infty$, we obtain alternative proof of the performance ratio $\frac{1}{2}(1 + \sqrt{2})$ by Kawaguchi and Kyan [14], avoiding the somewhat complicated transformation and case distinction in the proof of Lemma 10 and Schwegelshohn's proof [20]. For increasing m the tight performance ratio converges quite quickly to $\frac{1}{2}(1 + \sqrt{2})$: the difference lies in $O(1/m^2)$. By plugging in the machine-dependent performance ratio into Theorem 1, we obtain the following performance ratio for the WSEPT rule.

► **Corollary 15.** For instances with m machines of the problem $P|p_j \sim \text{stoch}|E[\sum w_j C_j]$ the WSEPT rule has performance ratio

$$1 + \frac{1}{2} \cdot \frac{\sqrt{(2m - k_m)k_m} - k_m}{m} \cdot (1 + \Delta).$$

This bound is better than the bound of Corollary 8 only if m and Δ both are small. Even for two machines, it is outdone for large Δ .

References

- 1 Adi Avidor, Yossi Azar, and Jiri Sgall. Ancient and new algorithms for load balancing in the l_p norm. *Algorithmica*, 29(3):422–441, 2001. doi:10.1007/s004530010051.
- 2 Chandra Chekuri, Rajeev Motwani, Balas Natarajan, and Clifford Stein. Approximation techniques for average completion time scheduling. *SIAM Journal on Computing*, 31(1):146–166, 2001. doi:10.1137/s0097539797327180.
- 3 Wang Chi Cheung, Felix Fischer, Jannik Matuschke, and Nicole Megow. A $\Omega(\Delta^{1/2})$ gap example for the WSEPT policy. Cited as personal communication in Uetz: MDS Autumn School Approximation Algorithms for Stochastic Optimization, 2014. URL: <http://www3.math.tu-berlin.de/MDS/summerschool14-material/Uetz-Exercises.pdf>.
- 4 Willard L. Eastman, Shimon Even, and I. Martin Isaacs. Bounds for the optimal scheduling of n jobs on m processors. *Management Science*, 11(2):268–279, 1964. doi:10.1287/mnsc.11.2.268.
- 5 Michel X. Goemans. Improved approximation algorithms for scheduling with release dates. In Michael Saks, editor, *Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 591–598. Society for Industrial and Applied Mathematics, 1997. URL: <http://dl.acm.org/citation.cfm?id=314394>.
- 6 Michel X. Goemans, Maurice Queyranne, Andreas S. Schulz, Martin Skutella, and Yao-guang Wang. Single machine scheduling with release dates. *SIAM J. Discrete Math.*, 15(2):165–192, 2002. doi:10.1137/S089548019936223X.
- 7 Ronald L. Graham, Eugene L. Lawler, Jan Karel Lenstra, and Alexander H. G. Rinnooy Kan. Optimization and approximation in deterministic sequencing and scheduling: a survey. In Peter L. Hammer, Ellis L. Johnson, and Bernhard H. Korte, editors, *Discrete Optimization II*, volume 5 of *Annals of Discrete Mathematics*, pages 287–326. Elsevier, 1979. doi:10.1016/S0167-5060(08)70356-X.
- 8 Varun Gupta, Benjamin Moseley, Marc Uetz, and Qiaomin Xie. Stochastic online scheduling on unrelated machines. In Friedrich Eisenbrand and Jochen Könemann, editors, *Integer Programming and Combinatorial Optimization - 19th International Conference, IPCO 2017, Waterloo, ON, Canada, June 26-28, 2017, Proceedings*, volume 10328 of *Lecture Notes in Computer Science*, pages 228–240. Springer, 2017. doi:10.1007/978-3-319-59250-3_19.
- 9 Leslie A. Hall, David B. Shmoys, and Joel Wein. Scheduling to minimize average completion time: Off-line and on-line algorithms. In Éva Tardos, editor, *Proceedings of the Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 142–151, Philadelphia, PA, USA, 1996. Society for Industrial and Applied Mathematics. URL: <http://dl.acm.org/citation.cfm?id=313852.313907>.
- 10 Sungjin Im, Benjamin Moseley, and Kirk Pruhs. Stochastic scheduling of heavy-tailed jobs. In Ernst W. Mayr and Nicolas Ollinger, editors, *32nd International Symposium on Theoretical Aspects of Computer Science, STACS 2015, March 4-7, 2015, Garching, Germany*, volume 30 of *LIPICs*, pages 474–486. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015. doi:10.4230/LIPICs.STACS.2015.474.
- 11 S. Jäger and M. Skutella. Generalizing the Kawaguchi-Kyan bound to stochastic parallel machine scheduling. *ArXiv e-prints*, jan 2018. arXiv:1801.01105.

- 12 Caroline Jagtenberg, Uwe Schwiegelshohn, and Marc Uetz. Analysis of smith's rule in stochastic machine scheduling. *Oper. Res. Lett.*, 41(6):570–575, 2013. doi:10.1016/j.orl.2013.08.001.
- 13 Christos Kalaitzis, Ola Svensson, and Jakub Tarnawski. Unrelated machine scheduling of jobs with uniform Smith ratios. In Philip N. Klein, editor, *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2654–2669, Philadelphia, PA, USA, 2017. Society for Industrial and Applied Mathematics. arXiv:1607.07631.
- 14 Tsuyoshi Kawaguchi and Seiki Kyan. Worst case bound of an LRF schedule for the mean weighted flow-time problem. *SIAM J. Comput.*, 15(4):1119–1129, 1986. doi:10.1137/0215081.
- 15 Nicole Megow, Marc Uetz, and Tjark Vredeveld. Models and algorithms for stochastic online scheduling. *Math. Oper. Res.*, 31(3):513–525, 2006. doi:10.1287/moor.1060.0201.
- 16 Rolf H. Möhring, Franz Josef Radermacher, and Gideon Weiss. Stochastic scheduling problems I - general strategies. *Zeitschr. für OR*, 28(7):193–260, 1984. doi:10.1007/BF01919323.
- 17 Rolf H. Möhring, Andreas S. Schulz, and Marc Uetz. Approximation in stochastic scheduling: the power of lp-based priority policies. *J. ACM*, 46(6):924–942, 1999. doi:10.1145/331524.331530.
- 18 Michael H. Rothkopf. Scheduling with random service times. *Management Science*, 12(9):707–713, may 1966. doi:10.1287/mnsc.12.9.707.
- 19 Andreas S. Schulz. Stochastic online scheduling revisited. In Boting Yang, Ding-Zhu Du, and Cao An Wang, editors, *Combinatorial Optimization and Applications, Second International Conference, COCOA 2008, St. John's, NL, Canada, August 21-24, 2008. Proceedings*, volume 5165 of *Lecture Notes in Computer Science*, pages 448–457. Springer, 2008. doi:10.1007/978-3-540-85097-7_42.
- 20 Uwe Schwiegelshohn. An alternative proof of the kawaguchi-kyan bound for the largest-ratio-first rule. *Oper. Res. Lett.*, 39(4):255–259, 2011. doi:10.1016/j.orl.2011.06.007.
- 21 Martin Skutella. A 2.542-approximation for precedence constrained single machine scheduling with release dates and total weighted completion time objective. *Oper. Res. Lett.*, 44(5):676–679, 2016. doi:10.1016/j.orl.2016.07.016.
- 22 Martin Skutella, Maxim Sviridenko, and Marc Uetz. Unrelated machine scheduling with stochastic processing times. *Math. Oper. Res.*, 41(3):851–864, 2016. doi:10.1287/moor.2015.0757.
- 23 Martin Skutella and Gerhard J. Woeginger. A PTAS for minimizing the total weighted completion time on identical parallel machines. *Math. Oper. Res.*, 25(1):63–75, 2000. doi:10.1287/moor.25.1.63.15212.
- 24 Wayne E. Smith. Various optimizers for single-stage production. *Naval Research Logistics Quarterly*, 3(1-2):59–66, mar 1956. doi:10.1002/nav.3800030106.
- 25 Richard R. Weber, Pravin Varaiya, and Jean Walrand. Scheduling jobs with stochastically ordered processing times on parallel machines to minimize expected flowtime. *Journal of Applied Probability*, 23(3):841–847, sep 1986. doi:10.2307/3214023.
- 26 Gideon Weiss. Approximation results in parallel machines stochastic scheduling. *Annals of Operations Research*, 26(1):195–242, 1990. doi:10.1007/BF02248591.
- 27 Gideon Weiss. Turnpike optimality of smith's rule in parallel machines stochastic scheduling. *Math. Oper. Res.*, 17(2):255–270, 1992. doi:10.1287/moor.17.2.255.