

A New Approach to Online Scheduling: Approximating the Optimal Competitive Ratio

Elisabeth Günther^{*†} Olaf Maurer^{*†} Nicole Megow^{*‡} Andreas Wiese^{§¶}

Abstract

We propose a new approach to competitive analysis in online scheduling by introducing the novel concept of competitive-ratio approximation schemes. Such a scheme algorithmically constructs an online algorithm with a competitive ratio arbitrarily close to the best possible competitive ratio for *any* online algorithm. We study the problem of scheduling jobs online to minimize the weighted sum of completion times on parallel, related, and unrelated machines, and we derive both deterministic and randomized algorithms which are almost best possible among all online algorithms of the respective settings. We also generalize our techniques to arbitrary monomial cost functions and apply them to the makespan objective. Our method relies on an abstract characterization of online algorithms combined with various simplifications and transformations. We also contribute algorithmic means to compute the actual value of the best possible competitive ratio up to an arbitrary accuracy. This strongly contrasts (nearly) all previous manually obtained competitiveness results and, most importantly, it reduces the search for the optimal competitive ratio to a question that a computer can answer. We believe that our concept can also be applied to many other problems and yields a new perspective on online algorithms in general.

1 Introduction

Competitive analysis [21,39] is the most popular method for studying the performance of online algorithms. It provides an effective framework to analyze and classify algorithms based on their worst-case behavior compared to an optimal offline algorithm over an infinite set of input instances. For some problem types, e.g., online

paging, competitive analysis may not be adequate to evaluate the performance of algorithms, but for a vast majority of online problems it is practical, natural, and yields meaningful results. A classical such problem is online scheduling to minimize the weighted average completion time. It has received a lot of attention in the past two decades. For different machine environments, a long sequence of papers emerged introducing new techniques and algorithms, improving upper and lower bounds on the competitive ratio of particular algorithms as well as on the best possible competitive ratio that any online algorithm can achieve. Still, unsatisfactory gaps remain. As for most online problems, a provably optimal online algorithm, w.r.t. competitive analysis, among *all* online algorithms is only known for very special cases.

In this work we close these gaps and present nearly optimal online scheduling algorithms. We provide *competitive-ratio approximation schemes* that compute algorithms with a competitive ratio that is at most a factor $1 + \varepsilon$ larger than the optimal ratio for any $\varepsilon > 0$. To that end, we introduce a new way of designing online algorithms. Apart from structuring and simplifying input instances, we find an abstract description of online scheduling algorithms, which allows us to reduce the infinite-size set of all online algorithms to a relevant set of finite size. This is the key for eventually allowing an enumeration scheme that finds an online algorithm with a competitive ratio arbitrarily close to the optimal one.

Besides improving on previous algorithms, our method also provides an algorithm to compute the competitive ratio of the designed algorithm, and even the best possible competitive ratio, up to any desired accuracy. This is clearly in strong contrast to all previously given (lower) bounds that stem from *manually* designed input instances. We are aware of only very few online problems for which a competitive ratio, or even the optimal competitive ratio, are known to be *computable* by some algorithm (for a not inherently finite problem). Our result is surprising, as there are typically no means of enumerating all possible input instances and all possible online algorithms. Even for only one given algorithm, usually one cannot compute its competitive ratio

^{*}Department of Mathematics, Technische Universität Berlin, Germany. Email: {eguenth,maurer,nmegow}@math.tu-berlin.de.

[†]Supported by the DFG Research center MATHEON *Mathematics for key technologies* in Berlin.

[‡]Supported by the German Science Foundation (DFG) under contract ME 3825/1.

[§]Department of Computer and System Sciences, Sapienza University of Rome, Italy. Email: wiese@dis.uniroma1.it.

[¶]Supported by the DFG Focus Program 1307 and by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD).

simply due to difficulties like the halting problem. We overcome these issues and pave the way for computer-assisted design of online algorithms.

We believe that our concept of abstraction for online algorithms can be applied successfully to other problems. We show this for other scheduling problems with jobs arriving online over time. We hope that our new approach to competitive analysis contributes to a better understanding of online algorithms and may lead to a new line of research in online optimization.

1.1 Problem Definition and Previous Results.

Competitive analysis. Given a minimization problem, a deterministic online algorithm A is called ρ -competitive if, for any problem instance \mathcal{I} , it achieves a solution of value $A(\mathcal{I}) \leq \rho \cdot \text{OPT}(\mathcal{I})$, where $\text{OPT}(\mathcal{I})$ denotes the value of an optimal offline solution for the same instance \mathcal{I} . A randomized online algorithm is called ρ -competitive, if it achieves in expectation a solution of value $\mathbb{E}[A(\mathcal{I})] \leq \rho \cdot \text{OPT}(\mathcal{I})$ for any instance \mathcal{I} . The *competitive ratio* ρ_A of A is the infimum over all ρ such that A is ρ -competitive. The minimum competitive ratio ρ^* achievable by any online algorithm is called *optimal*.

We define a competitive-ratio approximation scheme as a procedure that computes a nearly optimal online algorithm and at the same time provides a nearly exact estimate of the optimal competitive ratio.

DEFINITION 1.1. *A competitive-ratio approximation scheme computes for a given $\varepsilon > 0$ an online algorithm A with a competitive ratio $\rho_A \leq (1 + \varepsilon)\rho^*$. Moreover, it determines a value ρ' such that $\rho' \leq \rho^* \leq (1 + \varepsilon)\rho'$.*

Online scheduling. A scheduling instance consists of a fixed set of m machines and a set of jobs J , where each job $j \in J$ has processing time $p_j \in \mathbb{Q}^+$, weight $w_j \in \mathbb{Q}^+$, and release date $r_j \in \mathbb{Q}^+$. The jobs arrive online over time, i.e., each job becomes known to the scheduling algorithm only at its release date. We consider three different machine environments: identical parallel machines (denoted by P), related machines (Q) where each machine i has associated a speed s_i and processing a job j on machine i takes p_j/s_i time, and unrelated machines (R) where the processing time of a job j on each machine i is explicitly given as a value p_{ij} . The main problem considered in this paper is to schedule the jobs on the given set of machines so as to minimize $\sum_{j \in J} w_j C_j$, where C_j denotes the completion time of job j . We consider the problem with and without preemption. Using standard scheduling notation [15], we denote the non-preemptive (preemptive) problems that we consider in this paper by $\text{Pm}|r_j, (pmtn)|\sum w_j C_j$,

$\text{Qm}|r_j, (pmtn)|\sum w_j C_j$, and $\text{Rm}|r_j, pmtn|\sum w_j C_j$. We also briefly consider more general min-sum objectives $\sum_{j \in J} w_j f(C_j)$, where f is an arbitrary monomial function $f(x) = k \cdot x^\alpha$, with constant $\alpha \geq 1, k > 0$, and the classical makespan $C_{\max} := \max_{j \in J} C_j$.

Previous results. The offline variants of nearly all problems under consideration are NP-hard [12, 23, 24], but in most cases polynomial-time approximation schemes have been developed [1, 18, 19, 25].

The corresponding online settings have been a highly active field of research in the past fifteen years. A whole sequence of papers appeared introducing new algorithms, new relaxations and analytical techniques that decreased the gaps between lower and upper bounds on the optimal competitive ratio [2, 4–8, 11, 13, 14, 17, 20, 26, 27, 30–32, 34–38, 41]. Interestingly, despite the considerable effort, optimal competitive ratios are known only for $1|r_j, pmtn|\sum C_j$ [33] and for non-preemptive single-machine scheduling [2, 5, 20, 41]. In all other scheduling settings remain unsatisfactory, even quite significant gaps. We refer to our full version [16] for a detailed description of the state of the art for each individual problem.

Very recently, our new concept of competitive-ratio approximation schemes was applied also to job shop scheduling $\text{Jm}|r_j, op \leq \mu|C_{\max}$ and non-preemptive scheduling on unrelated machines $\text{Rm}|r_j|C_{\max}$ [22].

To the best of our knowledge, there are only very few problems in online optimization for which an optimal competitive ratio can be determined, bounded, or approximated by computational means. Lund and Reinhold [28] present a framework for upper-bounding the optimal competitive ratio of randomized algorithms by a linear program. For certain cases, e.g., the 2-server problem in a space of three points, this yields a provably optimal competitive ratio. Ebenlendr et al. [9, 10] study various online and semi-online variants of scheduling preemptive jobs on uniformly related machines to minimize the makespan. In contrast to our model, they assume the jobs to be given one by one (rather than over time). They prove that the optimal competitive ratio can be computed by a linear program for any given set of speeds. In terms of approximating the best possible performance guarantee, the work by Augustine, Irani, and Swamy [3] is closest to ours. They show how to compute a nearly optimal power-down strategy for a processor with a finite number of power states.

1.2 New Results and Methodology. In this paper, we introduce the concept of competitive-ratio approximation schemes and present such schemes for various scheduling problems with jobs arriving online over time. We present our technique fo-

cussing on the problems $\text{Pm}|r_j, (pmtn)|\sum w_j C_j$, $\text{Qm}|r_j, (pmtn)|\sum w_j C_j$ (assuming a constant range of machine speeds without preemption), and $\text{Rm}|r_j, pmtn|\sum w_j C_j$, and we comment on how it applies to other cost functions such as the makespan, C_{\max} , and $\sum_{j \in J} w_j f(C_j)$, where f is an arbitrary monomial function with fixed exponent. For any $\varepsilon > 0$, we show that the competitive ratios of our new algorithms are by at most a factor $1 + \varepsilon$ larger than the respective optimal competitive ratios. We obtain such nearly optimal online algorithms for the deterministic as well as the randomized setting, for any number of machines m . Moreover, we give an algorithm which estimates the optimal competitive ratio for these problems to any desired accuracy. Thus, we reduce algorithmically the performance gaps for all considered problems to an arbitrarily small value. These results reduce the long-time ongoing search for the best possible competitive ratio for the considered problems to a question that can be answered by a finite algorithm.

To achieve our results, we introduce a new and unusual way of designing online scheduling algorithms. We present an abstraction in which online algorithms are formalized as *algorithm maps*. Such a map receives as input a set of unfinished jobs together with the schedule computed so far. Based on this information, it returns a schedule for the next time instant. This view captures exactly how online algorithms operate under limited information. The total number of algorithm maps is unbounded. However, we show that there is a finite subset which approximates the entire set. More precisely, for any algorithm map there is a map in our subset whose competitive ratio is at most by a factor $1 + \varepsilon$ larger. To achieve this reduction, we first apply several standard techniques, such as geometric rounding, time-stretch, and weight-shift, to transform and simplify the input problem without increasing the objective value too much; see, e.g., [1]. The key, however, is the insight that it suffices for an online algorithm to base its decisions on the currently unfinished jobs and a very *limited part* of the so far computed schedule—rather than the entire history. This allows for an enumeration of all relevant algorithm maps (see also [29] for an enumeration routine for online algorithms for a fixed task system with finitely many states). For randomized algorithms we even show that we can restrict to instances with only constantly many jobs. As all our structural insights also apply to offline algorithms for the same problems, they might turn out to be useful for other settings as well.

Our algorithmic scheme contributes more than an improved competitive ratio. It also outputs (up to a factor $1 + \varepsilon$) the exact value of the competitive ratio of the

derived algorithm, which implies a $(1 + \varepsilon)$ -estimate for the optimal competitive ratio. This contrasts strongly all earlier results where (matching) upper and lower bounds on the competitive ratio of a particular and of all online algorithm had to be derived *manually*, instead of executing an algorithm using, e.g., a computer. In general, there are no computational means to determine the competitive ratio of an algorithm—even when it is a constant. It is simply not possible to enumerate all possible input instances. Even more, there are no general means of enumerating all possible online algorithms to determine the optimal competitive ratio. However, for the scheduling problems studied in this paper our extensive simplification of input instances and our abstract view on online algorithms allow us to overcome these obstacles, losing only a factor of $1 + \varepsilon$ in the objective.

Although the enumeration scheme for identifying the (nearly) optimal online algorithm heavily exploits unbounded computational resources, the resulting algorithm itself has polynomial running time. As a consequence, there are efficient online algorithms for the considered problems with almost optimal competitive ratios. Hence, the granted additional, even unbounded, computational power of online algorithms does not yield any significant benefit here.

Outline of the paper¹. In Section 2 we introduce several general transformations and observations that simplify the structural complexity of online scheduling in the setting of $\text{Pm}|r_j, pmtn|\sum w_j C_j$. Based on this, we present our abstraction of online algorithms and develop a competitive-ratio approximation scheme in Section 3. Next, we sketch in Section 4 how to extend these techniques to the non-preemptive setting and more general machine environments such that the approximation scheme (Sec. 3) remains applicable. In Section 5, we present competitive-ratio approximation schemes for the randomized setting. Finally, in Section 6, we extend our results to other objective functions.

2 General Simplifications and Techniques

In this section, we discuss several transformations that simplify the input and reduce the structural complexity of online schedules for $\text{Pm}|r_j, pmtn|\sum w_j C_j$. Later, we outline how to adapt these for more general settings. Our construction combines several transformation techniques known for offline PTASs (see [1] and the references therein) and a new technique to subdivide an instance online into parts which can be handled separately. We will use the terminology that *at* $1 + O(\varepsilon)$

¹Due to space constraints, we omit many proof details in this extended abstract. They can be found in our full version [16].

loss we can restrict to instances or schedules with certain properties. This means that we lose at most a factor $1+O(\varepsilon)$, as $\varepsilon \rightarrow 0$, by limiting our attention to those. We bound several relevant parameters by constants. If not stated differently, any mentioned constant depends only on ε and m .

LEMMA 2.1. ([1]) *At $1 + O(\varepsilon)$ loss we can restrict to instances where all processing times, release dates, and weights are powers of $1+\varepsilon$, no job is released before time $t = 1$, and $r_j \geq \varepsilon \cdot p_j$ for all jobs j .*

This standard *geometric rounding* procedure used in the lemma above allows us to see intervals of the form $I_x := [R_x, R_{x+1})$ with $R_x := (1+\varepsilon)^x$ as atomic entities. An online algorithm can define the corresponding schedule at the beginning of an interval since no further jobs are released until the next interval. Moreover, we assume at $1 + \varepsilon$ loss that all jobs that finish within I_x have completion time R_{x+1} .

Simplification within intervals. Our goal is to reduce the number of situations that can arise at the beginning of an interval. To this end, we partition the set of jobs released at time R_x into the set of *large* jobs L_x , with processing times at least $\varepsilon^3 R_x$, and the set of *small* jobs S_x with all remaining jobs. Running *Smith's Rule* [40] on small jobs allows us to group very small jobs to job *packs*, which we treat as single jobs. Together with Lemma 2.1 we obtain bounds on the lengths of jobs of each release date.

LEMMA 2.2. *At $1 + O(\varepsilon)$ loss, we can assume that for each interval I_x there are lower and upper bounds for the lengths of the jobs $S_x \cup L_x$ that are within a constant factor of R_x and the constants are independent of x . Also, the number of distinct processing times of jobs in each interval is upper-bounded by a constant.*

We look for jobs in S_x and L_x which can be excluded from processing within I_x at a loss of not more than $1 + O(\varepsilon)$. This allows us to bound the number of released jobs per interval.

LEMMA 2.3. *At $1 + O(\varepsilon)$ loss, we can restrict to instances where for each x , the number of jobs released at time R_x is bounded by a constant Δ .*

To prove the above lemmas, we use the technique of *time-stretching*, see [1]. In an online interpretation of this method, we shift the work assigned to any interval I_x to the interval I_{x+1} . This can be done at a loss of $1 + \varepsilon$ and we obtain free space of size $\varepsilon \cdot I_{x'-1}$ in each interval $I_{x'}$. Again using time-stretching, we can show that no job needs to be completed later than constantly many intervals after its release interval.

LEMMA 2.4. *There is a constant s such that at $1+O(\varepsilon)$ loss we can restrict to schedules such that for each interval I_x there is a subinterval of I_{x+s-1} which is large enough to process all jobs released at R_x and during which only those jobs are executed. We call this subinterval the safety net of interval I_x . We can assume that each job released at R_x finishes before time R_{x+s} .*

We can also simplify the complexity of the computed schedules by limiting the way jobs are preempted. We say that two large jobs are of the same *type* if they have the same processing time and the same release date. A job is *partially processed* if it has been processed, but not yet completed.

LEMMA 2.5. *There is a constant $\mu > 0$ such that at $1 + O(\varepsilon)$ loss we can restrict to schedules such that*

- *at the end of each interval, there are at most m large jobs of each type which are partially processed and each job j is processed to an extent which is a multiple of $p_j \cdot \mu$ and*
- *each small job finishes without preemption in the same interval where it started.*

Irrelevant history. The schedule that an online algorithm computes for an interval may depend on the set of currently unfinished jobs and possibly the entire schedule used so far. In the remainder of this section we show why we can assume that an online algorithm only takes a finite amount of history into account in its decision making, namely, the jobs with relatively large weight released in the last constantly many intervals.

Our strategy is to partition the time horizon into *periods*. For each integer $k \geq 0$, we define a period Q_k which consists of the s consecutive intervals $I_{k \cdot s}, \dots, I_{(k+1) \cdot s-1}$. For ease of notation, we will treat a period Q as the set of jobs released in that period. For a set of jobs J we denote by $rw(J) := \sum_{j \in J} r_j w_j$ their *release weight*. Note that $rw(J)$ forms a lower bound on the quantity that these jobs must contribute to the objective in *any* schedule. Due to Lemma 2.4, we also obtain an upper bound of $(1 + \varepsilon)^s \cdot rw(J)$ for the latter quantity.

LEMMA 2.6. *Let Q_k, \dots, Q_{k+p} be consecutive periods such that period Q_{k+p} is the first of this series with $rw(Q_{k+p}) \leq \frac{\varepsilon}{(1+\varepsilon)^s} \cdot \sum_{i=0}^{p-1} rw(Q_{k+i})$. Then at $1 + \varepsilon$ loss we can move all jobs in Q_{k+p} to their safety nets.*

The above observation defines a natural partition of a given instance \mathcal{I} into *parts* by the insignificant periods. Formally, let a_1, \dots, a_ℓ be all ordered indices such that Q_{a_i} is insignificant compared to the preceding

periods according to Lemma 2.6 ($a_0 := 0$). Let $a_{\ell+1}$ be the index of the last period. For each $i \in \{0, \dots, \ell\}$ we define a part P_i consisting of all periods $Q_{a_i+1}, \dots, Q_{a_{i+1}}$. Again, identify with P_i all jobs released in this part. We treat now each part P_i as a separate instance that we present to a given online algorithm. For the final output, we concatenate the computed schedules for the different parts. It then suffices to bound $A(P_i)/\text{OPT}(P_i)$ for each part P_i since $A(\mathcal{I})/\text{OPT}(\mathcal{I}) \leq \max_i \{A(P_i)/\text{OPT}(P_i)\} \cdot (1 + O(\varepsilon))$.

LEMMA 2.7. *At $1 + O(\varepsilon)$ loss we can restrict to instances which consist of only one part.*

Each but the last period of one part fulfills the opposite condition of the one from Lemma 2.6. This implies exponential growth for the series of partial sums of release weights (albeit with a small growth factor). From this observation, we get:

LEMMA 2.8. *There is a constant K such that the following holds: Let Q_1, Q_2, \dots, Q_p be consecutive periods such that $rw(Q_{i+1}) > \frac{\varepsilon}{(1+\varepsilon)^s} \cdot \sum_{\ell=1}^i rw(Q_\ell)$ for all i . Then in any schedule in which each job j finishes no later than by time $r_j \cdot (1 + \varepsilon)^s$ (e.g., using the safety net) it holds that $\sum_{i=1}^{p-K-1} \sum_{j \in Q_i} w_j C_j \leq \varepsilon \cdot \sum_{i=p-K}^p \sum_{j \in Q_i} w_j C_j$.*

The objective value of one part is therefore dominated by the contribution of the last K periods of this part. We will need this later to show that at $1 + \varepsilon$ loss we can assume that an online algorithm bases its decisions only on a constant amount of information. Denote the corresponding number of important intervals by $\Gamma := Ks$.

This enables us to partition the jobs into relevant and irrelevant jobs. Intuitively, a job is irrelevant if it is released very early (cf. Lemma 2.8) or its weight is very small in comparison to some other job. The subsequent lemma states that the irrelevant jobs can almost be ignored for the objective value of a schedule.

DEFINITION 2.1. *A job j is irrelevant at time R_x if it was irrelevant at time R_{x-1} , or $r_j < R_{x-\Gamma}$, or it is dominated at time R_x . This is the case if there is a job j' , either released at time R_x or already relevant at time R_{x-1} with release date at least $R_{x-\Gamma}$, such that $w_j < \frac{\varepsilon}{\Delta \cdot \Gamma \cdot (1+\varepsilon)^{\Gamma+s}} w_{j'}$. Otherwise, a job released until R_x is relevant at time R_x . Denote the respective subsets of some job set J by $\text{Rel}_x(J)$ and $\text{Ir}_x(J)$.*

LEMMA 2.9. *Consider a schedule of one part in which no job j finishes later than at time $r_j \cdot (1 + \varepsilon)^s$ (e.g., using the safety net) and let x be an interval index in this part. Then $\sum_{j \in \text{Ir}_x(J)} w_j C_j \leq O(\varepsilon) \cdot \sum_{j \in \text{Rel}_x(J)} w_j C_j$.*

The above lemma implies that at $1 + O(\varepsilon)$ loss we can restrict to online algorithms which schedule the remaining part of a job in its safety net, once it has become irrelevant.

3 Abstraction of Online Algorithms

In this section we show how to construct a competitive-ratio approximation scheme based on the observations of Section 2. To do so, we restrict ourselves to such simplified instances and schedules. The key idea is to characterize the behavior of an online algorithm by a map: For each interval, the map gets as input the schedule computed so far and all information about the currently unfinished jobs. Based on this information, the map outputs how to schedule the available jobs within this interval.

More precisely, we define the input by a *configuration* and the output by an *interval-schedule*.

DEFINITION 3.1. *An interval-schedule S for an interval I_x is defined by*

- *the index x of the interval,*
- *a set of jobs $J(S)$ available for processing in I_x together with the properties r_j, p_j, w_j of each job $j \in J(S)$ and its already finished part $f_j < p_j$ up to R_x ,*
- *for each job $j \in J(S)$ the information whether j is relevant at time R_x , and*
- *for each job $j \in J(S)$ and each machine i a value q_{ij} specifying for how long j is processed by S on machine i during I_x .*

An interval-schedule is called feasible if there is a feasible schedule in which the jobs of $J(S)$ are processed corresponding to the q_j values within the interval I_x . Denote the set of feasible interval-schedules as \mathcal{S} .

DEFINITION 3.2. *A configuration C for an interval I_x consists of*

- *the index x of the interval,*
- *a set of jobs $J(C)$ released up to time R_x together with the properties r_j, p_j, w_j, f_j of each job $j \in J(C)$,*
- *an interval-schedule for each interval $I_{x'}$ with $x' < x$.*

The set of all configurations is denoted by \mathcal{C} . An end-configuration is a configuration C for an interval I_x such that at time R_x , and not earlier, no jobs are left unprocessed.

We say that an interval-schedule S is *feasible* for a configuration C if the set of jobs in $J(C)$ which are unfinished at time R_x matches the set $J(S)$ with respect to release dates, total and remaining processing time, weight and relevance of the jobs.

Instead of online algorithms we work from now on with *algorithm maps*, which are defined as functions $f : \mathcal{C} \rightarrow \mathcal{S}$. An algorithm map determines a schedule $f(\mathcal{I})$ for a given scheduling instance \mathcal{I} by iteratively applying f to the corresponding configurations. W.l.o.g. we consider only algorithm maps f such that $f(C)$ is feasible for each configuration C and $f(\mathcal{I})$ is feasible for each instance \mathcal{I} . Like for online algorithms, we define the competitive ratio ρ_f of an algorithm map f by $\rho_f := \max_{\mathcal{I}} f(\mathcal{I})/\text{OPT}(\mathcal{I})$. Due to the following observation, algorithm maps are a natural generalization of online algorithms.

PROPOSITION 3.1. *For each online algorithm A there is an algorithm map f_A such that when A is in configuration $C \in \mathcal{C}$ at the beginning of an interval I_x , algorithm A schedules the jobs according to $f_A(C)$.*

Recall, that we restrict our attention to algorithm maps describing online algorithms which obey the simplifications introduced in Section 2. The essence of such online algorithms are the decisions for the relevant jobs. To this end, we define equivalence classes for configurations and for interval-schedules. Intuitively, two interval-schedules (configurations) are equivalent if we can obtain one from the other by scalar multiplication with the same value, while ignoring the irrelevant jobs.

DEFINITION 3.3. *Let S, S' be two feasible interval-schedules for two intervals $I_x, I_{x'}$. Denote by $J_{\text{Rel}}(S) \subseteq J(S)$ and $J_{\text{Rel}}(S') \subseteq J(S')$ the relevant jobs in $J(S)$ and $J(S')$. Let further $\sigma : J_{\text{Rel}}(S) \rightarrow J_{\text{Rel}}(S')$ be a bijection and y an integer. The interval-schedules S, S' are (σ, y) -equivalent if $r_{\sigma(j)} = r_j(1 + \varepsilon)^{x'-x}$, $p_{\sigma(j)} = p_j(1 + \varepsilon)^{x'-x}$, $f_{\sigma(j)} = f_j(1 + \varepsilon)^{x'-x}$, $q_{\sigma(j)} = q_j(1 + \varepsilon)^{x'-x}$ and $w_{\sigma(j)} = w_j(1 + \varepsilon)^y$ for all $j \in J_{\text{Rel}}(S)$. The interval-schedules S, S' are equivalent (denoted by $S \sim S'$) if a map σ and an integer y exist such that they are (σ, y) -equivalent.*

DEFINITION 3.4. *Let C, C' be two configurations for two intervals $I_x, I_{x'}$. Denote by $J_{\text{Rel}}(C), J_{\text{Rel}}(C')$ the jobs which are relevant at times $R_x, R_{x'}$ in C, C' , respectively. Configurations C, C' are equivalent (denoted by $C \sim C'$) if there is a bijection $\sigma : J_{\text{Rel}}(C) \rightarrow J_{\text{Rel}}(C')$ and an integer y such that*

- $r_{\sigma(j)} = r_j(1 + \varepsilon)^{x'-x}$, $p_{\sigma(j)} = p_j(1 + \varepsilon)^{x'-x}$, $f_{\sigma(j)} = f_j(1 + \varepsilon)^{x'-x}$ and $w_{\sigma(j)} = w_j(1 + \varepsilon)^y$ for all $j \in J_{\text{Rel}}(C)$, and

- the interval-schedules of I_{x-k} and $I_{x'-k}$ are (σ, y) -equivalent when restricted to the jobs in $J_{\text{Rel}}(C)$ and $J_{\text{Rel}}(C')$ for each $k \in \mathbb{N}$.

A configuration C is *realistic* for an algorithm map f if there is an instance \mathcal{I} such that if f processes \mathcal{I} then at time R_x it is in configuration C . The following lemma shows that we can restrict the set of algorithm maps under consideration to those which treat equivalent configurations equivalently. We call algorithm maps obeying this condition in addition to the restrictions of Section 2 *simplified algorithm maps*.

LEMMA 3.1. *At $1 + O(\varepsilon)$ loss we can restrict to algorithm maps f such that $f(C) \sim f(C')$ for any two equivalent configurations C, C' .*

Proof. Let f be an algorithm map. For each equivalence class $\mathcal{C}' \subseteq \mathcal{C}$ of the set of configurations we pick a representative C which is realistic for f . For each configuration $C' \in \mathcal{C}'$, we define a new algorithm map \bar{f} by setting $\bar{f}(C')$ to be the interval-schedule for C' which is equivalent to $f(C)$. One can show by induction that \bar{f} is always in a configuration such that an equivalent configuration is realistic for \bar{f} . Hence, equivalence classes without realistic configurations for f are not relevant. We claim that $\rho_{\bar{f}} \leq (1 + O(\varepsilon))\rho_f$.

Consider an instance $\bar{\mathcal{I}}$. We show that there is an instance \mathcal{I} such that $\bar{f}(\bar{\mathcal{I}})/\text{OPT}(\bar{\mathcal{I}}) \leq (1 + O(\varepsilon))f(\mathcal{I})/\text{OPT}(\mathcal{I})$ which implies the claim. Consider an interval $I_{\bar{x}}$. Let \bar{C} be the end-configuration obtained when \bar{f} is applied iteratively on $\bar{\mathcal{I}}$. Let C be the representative of the equivalence class of \bar{C} , which was chosen above and which is realistic for f (C is also an end-configuration). Therefore, there is an instance \mathcal{I} such that C is reached at time R_x when f is applied on \mathcal{I} . Hence, \mathcal{I} is the required instance since the relevant jobs dominate the objective value (see Lemma 2.9) and $C \sim \bar{C}$. \square

LEMMA 3.2. *There are only constantly many simplified algorithm maps. Each simplified algorithm map can be described using finite information.*

Proof. From the simplifications introduced in Section 2 follows that the domain of the algorithm maps under consideration contains only constantly many equivalence classes of configurations. Also, the target space contains only constantly many equivalence classes of interval-schedules. For an algorithm map f which obeys the restrictions of Section 2, the interval-schedule $f(C)$ is fully specified when knowing only C and the equivalence class which contains $f(C)$ (since the irrelevant jobs are moved to their safety net anyway). Since $f(C) \sim f(C')$ for a simplified algorithm map f if $C \sim C'$, we

conclude that there are only constantly many simplified algorithm maps. Finally, each equivalence class of configurations and interval-schedules can be characterized using only finite information, and hence the same holds for each simplified algorithm map. \square

The next lemma shows that up to a factor $1 + \varepsilon$ worst case instances of simplified algorithm maps span only constantly many intervals. Using this property, we will show in the subsequent lemmas that the competitive ratio of a simplified algorithm map can be determined algorithmically up to a $1 + \varepsilon$ factor.

LEMMA 3.3. *There is a constant E such that for any instance I and any simplified algorithm map f there is a realistic end-configuration \tilde{C} for an interval $I_{\tilde{x}}$ with $\tilde{x} \leq E$ which is equivalent to the corresponding end-configuration when f is applied to I .*

Proof. Consider a simplified algorithm map f . For each interval I_x , denote by \mathcal{C}_x^f the set of realistic equivalence classes for I_x , i.e., the equivalence classes which have a realistic representative for I_x . Since there are constantly many equivalence classes and thus constantly many sets of equivalence classes, there must be a constant E independent of f such that $\mathcal{C}_{\tilde{x}}^f = \mathcal{C}_{\tilde{x}'}^f$ for some $\tilde{x} < \tilde{x}' \leq E$. Since f is simplified it can be shown by induction that $\mathcal{C}_{\tilde{x}+k}^f = \mathcal{C}_{\tilde{x}'+k}^f$ for any $k \in \mathbb{N}$, i.e., f cycles with period length $\tilde{x}' - \tilde{x}$.

Consider now some instance I and let C with interval I_x be the corresponding end-configuration when f is applied to I . If $x \leq E$ we are done. Otherwise there must be some $k \leq \tilde{x}' - \tilde{x}$ such that $\mathcal{C}_{\tilde{x}+k}^f = \mathcal{C}_x^f$ since f cycles with this period length. Hence, by definition of $\mathcal{C}_{\tilde{x}+k}^f$ there must be a realistic end-configuration \tilde{C} which is equivalent to C for the interval $I_{\tilde{x}}$ with $\tilde{x} := \tilde{x} + k \leq E$. \square

LEMMA 3.4. *Let f be a simplified algorithm map. There is an algorithm which approximates ρ_f up to a factor $1 + \varepsilon$, i.e., it computes a value ρ' with $\rho' \leq \rho_f \leq (1 + O(\varepsilon))\rho'$.*

Proof sketch. By Lemma 2.9, the relevant jobs in a configuration dominate the entire objective value. In particular, we do not need to know the irrelevant jobs of a configuration if we only want to approximate its objective value up to a factor of $1 + O(\varepsilon)$. For an end-configuration C denote by $val_C(J_{\text{Rel}}(C))$ the objective value of the jobs in $J_{\text{Rel}}(C)$ in the history of C . We define $r(C) := val_C(J_{\text{Rel}}(C))/\text{OPT}(J_{\text{Rel}}(C))$ to be the achieved competitive ratio of C when restricted to the relevant jobs. According to Lemma 3.3, it suffices to construct the sets $\mathcal{C}_0^f, \dots, \mathcal{C}_E^f$ in order to approximate

the competitive ratio of all end-configurations in these sets. We start with \mathcal{C}_0^f and determine $f(C)$ for one representant C of each equivalence class $\mathcal{C} \in \mathcal{C}_0^f$. Based on this we determine the set \mathcal{C}_1^f . We continue inductively to construct all sets \mathcal{C}_x^f with $x \leq E$.

We define r_{\max} to be the maximum ratio $r(C)$ for an end-configuration $C \in \cup_{0 \leq x \leq E} \mathcal{C}_x^f$. Due to Lemma 3.3 and Lemma 2.9 the value r_{\max} implies the required ρ' fulfilling the properties claimed in this lemma. \square

Our main algorithm works as follows. We first enumerate all simplified algorithm maps. For each simplified algorithm map f we approximate ρ_f using Lemma 3.4. We output the map f with the minimum (approximated) competitive ratio. Note that the resulting online algorithm has polynomial running time: All simplifications of a given instance can be done efficiently and for a given configuration, the equivalence class of the schedule for the next interval can be found in a look-up table of constant size.

THEOREM 3.1. *For any $m \in \mathbb{N}$ we obtain a competitive-ratio approximation scheme for $\text{Pm} | r_j, \text{pmtn} | \sum w_j C_j$.*

4 Extensions to Other Settings

4.1 Non-preemptive Scheduling. When preemption is not allowed, the definition of the safety net (Lemma 2.4) needs to be adjusted since we cannot ensure that at the end of each interval I_{x+s} there is a machine idle. However, we can guarantee that there is a reserved space *somewhere* in $[R_x, R_{x+s})$ to process the small and big jobs in $S_x \cup L_x$. Furthermore, we cannot enforce that a big job j is processed for exactly a certain multiple of $p_j \mu$ in each interval (Lemma 2.5). To solve this, we pretend that we could preempt j and ensure that after j has been preempted its machine stays idle until j continues. Next, we can no longer assume that each part can be treated independently (Lemma 2.7). Since some of the remaining jobs at the end of a part may have already started processing, we cannot simply move them to their safety net. Here we use the following simplification.

LEMMA 4.1. *Let $\text{first}(i)$ denote the job that is released first in part P_i . At $1 + \varepsilon$ loss, we can restrict to instances such that $\sum_{\ell=1}^{i-1} rw(P_\ell) \leq \frac{\varepsilon}{(1+\varepsilon)^8} \cdot rw(\text{first}(i))$, i.e., $\text{first}(i)$ dominates all previous parts.*

Therefore, at $1 + \varepsilon$ loss it is enough to consider only the currently running jobs from the previous part and the last Γ intervals from the current part when taking decisions. Finally, we add some minor modifications to handle the case that a currently running job is

dominated by some other job. With these adjustments, we have only constantly many equivalence classes for interval-schedules and configurations, which allows us to construct a competitive-ratio approximation scheme as in Section 3.

THEOREM 4.1. *For any $m \in \mathbb{N}$ we obtain a competitive-ratio approximation scheme for $\text{Pm} | r_j | \sum w_j C_j$.*

4.2 Scheduling on Related Machines. In this setting, each machine i has associated a speed s_i , such that processing job j on machine i takes p_j/s_i time units. W.l.o.g. the slowest machine has unit speed. Let s_{\max} denote the maximum speed in an instance. An adjusted version of Lemma 2.1 ensures that at $1 + \varepsilon$ loss $r_j \geq \varepsilon p_j/s_{\max}$ for all jobs j (rather than $r_j \geq \varepsilon p_j$). Furthermore, we can bound the number of distinct processing times and the number of released jobs of each interval, using similar arguments as in the unit-speed case.

LEMMA 4.2. *At $1 + O(\varepsilon)$ loss we can restrict to instances where for each release date the number of released jobs and the number of distinct processing times is bounded by a constant depending only on ε , m , and s_{\max} .*

We establish the safety net for the jobs of each release date R_x only on the *fastest* machine and thereby ensure the condition of Lemma 2.4 in the related machine setting. For the non-preemptive setting we incorporate the adjustments introduced in Section 4.1. Since at $1 + \varepsilon$ loss we can round the speeds of the machines to powers of $1 + \varepsilon$ we obtain the following result.

THEOREM 4.2. *For any $m \in \mathbb{N}$ we obtain competitive-ratio approximation schemes for $\text{Qm} | r_j, pmtn | \sum w_j C_j$ and $\text{Qm} | r_j | \sum w_j C_j$, assuming that the speeds of any two machines differ by at most a constant factor.*

In the preemptive setting we can strengthen the result and give a competitive-ratio approximation scheme for the case that machine speeds are part of the input, that is, we obtain a nearly optimal competitive ratio for *any* speed vector. The key is to bound the variety of different speeds. To that end, we show that at $1 + \varepsilon$ loss a very fast machine can simulate $m - 1$ very slow machines.

LEMMA 4.3. *For $\text{Qm} | r_j, pmtn | \sum w_j C_j$, at $1 + O(\varepsilon)$ loss, we can restrict to instances in which s_{\max} is bounded by m/ε .*

As speeds are geometrically rounded, we have for each value m only finitely many speed vectors. Thus, our

enumeration scheme finds a nearly optimal online algorithm with a particular routine for each speed vector.

THEOREM 4.3. *For any $m \in \mathbb{N}$ we obtain a competitive-ratio approximation scheme for $\text{Qm} | r_j, pmtn | \sum w_j C_j$.*

4.3 Preemptive Scheduling on Unrelated Machines. When each job j has its individual processing time p_{ij} on machine i , the problem complexity increases significantly. We restrict to preemptive scheduling and show how to decrease the complexity to apply our approximation scheme. The key is to bound the range of the finite processing times for each job (which is unfortunately not possible in the non-preemptive case, see [1] for a counterexample).

LEMMA 4.4. *At $1 + \varepsilon$ loss we can restrict to instances in which for each job j the ratio of any two of its finite processing times is bounded by m/ε .*

The above lemma allows us to introduce the notion of *job classes*. Two jobs j, j' are of the same class if they have finite processing times on exactly the same machines and $p_{ij}/p_{ij'} = p_{i'j}/p_{i'j'}$ for any two such machines i and i' . For fixed m , the number of different job classes is bounded by a constant W .

For each job class, we define large and small tasks similar to the identical machine case: for each job j we define a value $\tilde{p}_j := \max_i \{p_{ij} | p_{ij} < \infty\}$ and say a job is *large* if $\tilde{p}_j \geq \varepsilon^2 r_j/W$ and *small* otherwise. For each job class separately, we perform the adjustments of Section 2. This yields the following lemma.

LEMMA 4.5. *At $1 + O(\varepsilon)$ loss we can restrict to instances and schedules such that*

- *for each job class, the number of distinct values \tilde{p}_j of jobs j with the same release date is bounded by a constant,*
- *for each job class, the number of jobs with the same release date is bounded by a constant $\tilde{\Delta}$,*
- *a large job j is only preempted at integer multiples of $\tilde{p}_j \cdot \tilde{\mu}$ for some constant $\tilde{\mu}$ and small jobs are never preempted and finish in the same interval where they start.*

The above lemmas imply that both, the number of equivalence classes of configurations and the number of equivalence classes for interval-schedules are bounded by constants. Thus, we can apply the enumeration scheme from Section 3.

THEOREM 4.4. *For any $m \in \mathbb{N}$ we obtain a competitive-ratio approximation scheme for $\text{Rm} | r_j, pmtn | \sum w_j C_j$.*

5 Randomized algorithms

When algorithms are allowed to make random choices and we consider expected values of schedules, we can restrict to instances which span only constantly many periods. Assuming the simplifications of Section 2, this allows a restriction to instances with a constant number of jobs.

LEMMA 5.1. *For randomized algorithms, at $1 + \varepsilon$ loss we can restrict to instances in which all jobs are released in at most $(1 + \varepsilon)^s / \varepsilon$ consecutive periods.*

Proof idea. Beginning at a randomly chosen period Q_i with $i \in [0, M)$, with $M := \lceil (1 + \varepsilon)^s / \varepsilon \rceil$, we move all jobs released in Q_{i+kM} , $k = 0, 1, \dots$, to their safety net. At $1 + \varepsilon$ loss, this gives us a partition into parts, at the end of which no job remains, and we can treat each part independently. \square

A randomized online algorithm can be viewed as a function that maps every possible configuration C to a probability distribution of interval-schedules which are feasible for C . To apply our algorithmic framework from the deterministic setting that enumerates all algorithm maps, we discretize the probability space and define discretized algorithm maps. To this end, let $\bar{\Gamma}$ denote the maximum number of intervals in instances with at most $(1 + \varepsilon)^s / \varepsilon$ periods.

DEFINITION 5.1. (DISCRETIZED ALGORITHM MAPS)
Let $\bar{\mathcal{C}}$ be the set of configurations for intervals I_x with $x \leq \bar{\Gamma}$, let $\bar{\mathcal{S}}$ be the set of interval-schedules for intervals I_x with $x \leq \bar{\Gamma}$, and let $\delta > 0$. A δ -discretized algorithm map is a function $f : \bar{\mathcal{C}} \times \bar{\mathcal{S}} \rightarrow [0, 1]$ such that $f(C, S) = k \cdot \delta$ with some $k \in \mathbb{N}_0$ for all $C \in \bar{\mathcal{C}}$ and $S \in \bar{\mathcal{S}}$, and $\sum_{S \in \bar{\mathcal{S}}} f(C, S) = 1$ for all $C \in \bar{\mathcal{C}}$.

By restricting to δ -discretized algorithm maps we do not lose too much in the competitive ratio.

LEMMA 5.2. *There is a value $\delta > 0$ such that for any (randomized) algorithm map f there is a δ -discretized randomized algorithm map g with $\rho_g \leq \rho_f (1 + \varepsilon)$.*

Proof idea. Let f be a randomized algorithm map and let $\delta > 0$ such that $1/\delta \in \mathbb{N}$. We define a new δ -discretized algorithm map g . For each configuration C we define the values $g(C, S)$ such that $\lfloor f(C, S)/\delta \rfloor \cdot \delta \leq g(C, S) \leq \lceil f(C, S)/\delta \rceil \cdot \delta$ and $\sum_{S \in \bar{\mathcal{S}}} g(C, S) = 1$. To see that $\rho_g \leq (1 + \varepsilon)\rho_f$, consider an instance \mathcal{I} and a possible schedule $S(\mathcal{I})$ for \mathcal{I} . There is a probability p that f outputs $S(\mathcal{I})$. We show that the schedules which have large probability p dominate $\mathbb{E}[f(\mathcal{I})]/\text{OPT}(\mathcal{I})$. We show further that if p is sufficiently large, the

probability that g produces $S(\mathcal{I})$ is in $[p/(1 + \varepsilon), p(1 + \varepsilon))$, which implies the Lemma. \square

Like in the deterministic case, we can now show that at $1 + \varepsilon$ loss it suffices to restrict to *simplified δ -discretized algorithm maps* which treat equivalent configurations equivalently, similar to Lemma 3.1 (replacing Γ by $\bar{\Gamma}$ in Definition 2.1 of the irrelevant jobs). As there are only constantly many of these maps, we enumerate all of them, test each map for its competitive ratio, and select the best of them.

THEOREM 5.1. *We obtain randomized competitive-ratio approximation schemes for $\text{Pm}|r_j, (\text{pmtn})|\sum w_j C_j$, $\text{Qm}|r_j, \text{pmtn}|\sum w_j C_j$, and $\text{Rm}|r_j, \text{pmtn}|\sum w_j C_j$ and for $\text{Qm}|r_j|\sum w_j C_j$ with a bounded range of speeds for any fixed $m \in \mathbb{N}$.*

6 General Min-Sum Objectives and Makespan

In this section we briefly argue how the techniques presented above for minimizing $\sum_j w_j C_j$ can be used for constructing online algorithm schemes for other scheduling problems with jobs arriving online over time, namely for minimizing $\sum_{j \in J} w_j f(C_j)$, with $f(x) = k \cdot x^\alpha$ and constant $\alpha \geq 1, k > 0$, and the makespan.

Since monomial functions f have the property that $f((1 + \varepsilon)C_j) \leq (1 + O(\varepsilon))f(C_j)$, the arguments in previous sections apply almost directly to the generalized min-sum objective. In each step of simplification and abstraction, we have an increased loss in the performance guarantee, but it is covered by the $O(\varepsilon)$ -term.

Consider now the makespan objective. The *simplifications within intervals* of Section 2 are based on arguing on completion times of individual jobs, and clearly hold also for the last job. Thus, they directly apply to makespan minimization. We simplify the definition of *irrelevant history* by omitting the partition of the instance into parts and we define a job j to be *irrelevant at time R_x* if $r_j \leq R_{x-s}$.

Based on this definition, we define equivalence classes of configurations (ignoring weights and previous interval-schedules) and again restrict to algorithm maps f with $f(C) \sim f(C')$ for any two equivalent configurations (Lemma 3.6). Lemmas 3.7–3.9 then hold accordingly and yield a competitive-ratio approximation scheme. Finally, the adjustments of Sections 4 and 5 can be made accordingly. Without the partition of the instance into parts, this even becomes easier in the non-preemptive setting. Thus, we can state the following result.

THEOREM 6.1. *For any $m \in \mathbb{N}$ there are deterministic and randomized competitive-ratio approximation schemes for preemptive and non-preemptive scheduling,*

on m identical, related (with bounded speed ratio when non-preemptive), and unrelated machines (only preemptive) for the objectives of minimizing C_{\max} and minimizing $\sum_{j \in J} w_j f(C_j)$, with $f(x) = k \cdot x^\alpha$ and constant $\alpha \geq 1, k > 0$.

7 Conclusions

We introduce the concept of competitive-ratio approximation schemes that compute online algorithms with a competitive ratio arbitrarily close to the best possible competitive ratio. We provide such schemes for various problem variants of scheduling jobs online to minimize the weighted sum of completion times, arbitrary monomial cost functions, and the makespan.

The techniques derived in this paper provide a new and interesting view on the behavior of online algorithms. We believe that they contribute to the understanding of such algorithms and possibly open a new line of research in which they yield even further insights. In particular, it seems promising that our methods could also be applied to other online problems than scheduling jobs arriving online over time.

References

- [1] F. N. Afrati, E. Bampis, C. Chekuri, D. R. Karger, C. Kenyon, S. Khanna, I. Milis, M. Queyranne, M. Skutella, C. Stein, and M. Sviridenko. Approximation schemes for minimizing average weighted completion time with release dates. In *Proceedings of the 40th IEEE Symposium on the Foundations of Computer Science (FOCS)*, pages 32–43, 1999.
- [2] E. J. Anderson and C. N. Potts. On-line scheduling of a single machine to minimize total weighted completion time. *Mathematics of Operations Research*, 29:686–697, 2004.
- [3] J. Augustine, S. Irani, and C. Swamy. Optimal power-down strategies. *SIAM Journal on Computing*, 37(5):1499–1516, Jan. 2008.
- [4] S. Chakrabarti, C. A. Phillips, A. S. Schulz, D. B. Shmoys, C. Stein, and J. Wein. Improved algorithms for minsum criteria. In *Proceedings of the 23rd International Conference on Automata, Languages and Programming (ICALP)*, volume 1099, pages 646–657, 1996.
- [5] C. Chekuri, R. Motwani, B. Natarajan, and C. Stein. Approximation techniques for average completion time scheduling. *SIAM Journal on Computing*, 31:146–166, 2001.
- [6] B. Chen and A. P. Vestjens. Scheduling on identical machines: How good is LPT in an on-line setting? *Operations Research Letters*, 21:165–169, 1997.
- [7] C. Chung, T. Nonner, and A. Souza. SRPT is 1.86-competitive for completion time scheduling. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1373–1388. SIAM, 2010.
- [8] J. Correa and M. Wagner. LP-based online scheduling: From single to parallel machines. *Mathematical Programming*, 119:109–136, 2009.
- [9] T. Ebenlendr, W. Jawor, and J. Sgall. Preemptive online scheduling: Optimal algorithms for all speeds. *Algorithmica*, 53:504–522, 2009.
- [10] T. Ebenlendr and J. Sgall. Semi-online preemptive scheduling: One algorithm for all variants. *Theory Comput. Syst.*, 48(3):577–613, 2011.
- [11] L. Epstein and R. van Stee. Lower bounds for on-line single-machine scheduling. *Theoretical Computer Science*, 299:439–450, 2003.
- [12] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., 1979.
- [13] M. X. Goemans. Improved approximation algorithms for scheduling with release dates. In *Proceedings of the 8th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 591–598. SIAM, 1997.
- [14] M. X. Goemans, M. Queyranne, A. S. Schulz, M. Skutella, and Y. Wang. Single machine scheduling with release dates. *SIAM Journal on Discrete Mathematics*, 15:165–192, 2002.
- [15] R. L. Graham, E. L. Lawler, J. K. Lenstra, and A. H. G. Rinnooy Kan. Optimization and approximation in deterministic sequencing and scheduling: A survey. *Annals of Discrete Mathematics*, 5:287–326, 1979.
- [16] E. Günther, O. Maurer, N. Megow, and A. Wiese. A new approach to online scheduling: Approximating the optimal competitive ratio. Technical report, CORR, abs/1204.0897, <http://arxiv.org/abs/1204.0897>, 2012.
- [17] L. A. Hall, A. S. Schulz, D. B. Shmoys, and J. Wein. Scheduling to minimize average completion time: Off-line and on-line approximation algorithms. *Mathematics of Operations Research*, 22:513–544, 1997.
- [18] D. S. Hochbaum and D. B. Shmoys. Using dual approximation algorithms for scheduling problems theoretical and practical results. *J. ACM*, 34:144–162, 1987.
- [19] D. S. Hochbaum and D. B. Shmoys. A polynomial approximation scheme for scheduling on uniform processors: Using the dual approximation approach. *SIAM Journal on Computing*, 17:539–551, 1988.
- [20] H. Hoogeveen and A. P. A. Vestjens. Optimal on-line algorithms for single-machine scheduling. In W. H. Cunningham, S. T. McCormick, and M. Queyranne, editors, *Proceedings of the 5th International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, volume 1084, pages 404–414, 1996.
- [21] A. R. Karlin, M. S. Manasse, L. Rudolph, and D. D. Sleator. Competitive snoopy paging. *Algorithmica*, 3:70–119, 1988.
- [22] A. Kurpisz, M. Mastrolilli, and G. Stamoulis. Competitive ratio approximation schemes for makespan scheduling problems. In *Proceedings of the 10th*

- Workshop on Approximation and Online Algorithms (WAOA 2012)*, 2012. To appear.
- [23] J. Labetoulle, E. L. Lawler, J. K. Lenstra, and A. H. G. Rinnooy Kan. Preemptive scheduling of uniform machines subject to release dates. In *Progress in Combinatorial Optimization*, pages 245–261. Academic Press Canada, 1984.
- [24] J. K. Lenstra, A. H. G. Rinnooy Kan, and P. Brucker. Complexity of machine scheduling problems. *Annals of Discrete Mathematics*, 1:243–362, 1977.
- [25] J. K. Lenstra, D. B. Shmoys, and É. Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical Programming*, 46:259–271, 1990.
- [26] P. Liu and X. Lu. On-line scheduling of parallel machines to minimize total completion times. *Computers and Operations Research*, 36:2647–2652, 2009.
- [27] X. Lu, R. A. Sitters, and L. Stougie. A class of on-line scheduling algorithms to minimize total completion time. *Operations Research Letters*, 31:232–236, 2003.
- [28] C. Lund and N. Reingold. Linear programs for randomized on-line algorithms. In *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, SODA '94, pages 382–391, Philadelphia, PA, USA, 1994. Society for Industrial and Applied Mathematics.
- [29] M. Manasse, L. McGeoch, and D. Sleator. Competitive algorithms for on-line problems. In *Proceedings of the 20th annual ACM Symposium on Theory of Computing (STOC)*, pages 322–333. ACM, 1988.
- [30] N. Megow. *Coping with incomplete information in scheduling—stochastic and online models*. PhD thesis, Technische Universität Berlin, Germany, 2007.
- [31] N. Megow and A. S. Schulz. On-line scheduling to minimize average completion time revisited. *Operations Research Letters*, 32:485–490, 2004.
- [32] C. A. Phillips, C. Stein, and J. Wein. Minimizing average completion time in the presence of release dates. *Mathematical Programming*, 82:199–223, 1998.
- [33] L. Schrage. A proof of the optimality of the shortest remaining processing time discipline. *Operations Research*, 16:687–690, 1968.
- [34] A. S. Schulz and M. Skutella. The power of α -points in preemptive single machine scheduling. *Journal of Scheduling*, 5:121–133, 2002.
- [35] A. S. Schulz and M. Skutella. Scheduling unrelated machines by randomized rounding. *SIAM Journal on Discrete Mathematics*, 15:450–469, 2002.
- [36] S. S. Seiden. A guessing game and randomized online algorithms. In *Proceedings of the 32nd ACM Symposium on the Theory of Computing (STOC)*, pages 592–601, 2000.
- [37] R. Sitters. Competitive analysis of preemptive single-machine scheduling. *Operations Research Letters*, 38:585–588, 2010.
- [38] R. Sitters. Efficient algorithms for average completion time scheduling. In *Proceedings of the 14th International Conference on Integer Programming and Combinatorial Optimization, (IPCO)*, volume 6080, pages 411–423. Springer, 2010.
- [39] D. D. Sleator and R. E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28:202–208, 1985.
- [40] W. E. Smith. Various optimizers for single-stage production. *Naval Research Logistics Quarterly*, 3:59–66, 1956.
- [41] L. Stougie and A. P. A. Vestjens. Randomized on-line scheduling: How low can't you go? *Operations Research Letters*, 30:89–96, 2002.