



HAL
open science

Scheduling with a processing time oracle

Fanny Dufossé, Christoph Dürr, Noël Nadal, Denis Trystram, Óscar Vásquez

► **To cite this version:**

Fanny Dufossé, Christoph Dürr, Noël Nadal, Denis Trystram, Óscar Vásquez. Scheduling with a processing time oracle. Applied Mathematical Modelling, 2022, 104, pp.701-720. 10.1016/j.apm.2021.12.020 . hal-03523262

HAL Id: hal-03523262

<https://hal.science/hal-03523262v1>

Submitted on 8 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Scheduling with a processing time oracle

Fanny Dufossé^a, Christoph Dürr^{b,*}, Noël Nadal^c, Denis Trystram^a, Óscar C. Vásquez^{d,e}

^aUniv. Grenoble-Alpes, Inria, CNRS, France.

^bSorbonne University, CNRS, LIP6, France.

^cEcole Normale Supérieure Paris-Saclay, Gif-sur-Yvette, France.

^dUniversity of Santiago of Chile (USACH), Faculty of Engineering, Program for the Development of Sustainable Production Systems (PDSPS), Chile.

^eUniversity of Santiago of Chile (USACH), Faculty of Engineering, Industrial Engineering Department, Chile.

Abstract

In this paper we study a single machine scheduling problem with the objective of minimizing the sum of completion times. Each of the given jobs is either short or long. However the processing times are initially hidden to the algorithm, but can be tested. This is done by executing a processing time oracle, which reveals the processing time of a given job. Each test occupies a time unit in the schedule, therefore the algorithm must decide for which jobs it will call the processing time oracle. The objective value of the resulting schedule is compared with the objective value of an optimal schedule, which is computed using full information. The resulting competitive ratio measures the price of hidden processing times, and the goal is to design an algorithm with minimal competitive ratio.

Two models are studied in this paper. In the *non-adaptive* model, the algorithm needs to decide beforehand which jobs to test, and which jobs to execute untested. However in the *adaptive* model, the algorithm can make these decisions adaptively depending on the outcomes of the job tests. In both models we provide optimal polynomial time algorithms following a *two-phase strategy*, which consist of a first phase where jobs are tested, and a second phase where jobs are executed obliviously. Experiments give strong evidence that optimal algorithms have this structure. Proving this property is left as an open problem.

Keywords: scheduling, uncertainty, competitive ratio, processing time oracle

1. Introduction

A typical combinatorial optimization problem consists of a clear defined input, for which the algorithm has to compute a solution minimizing some cost. This is the beautiful simple world of theory. In contrast, everyone who participated in some industrial project can testify that obtaining the input is one of the hardest aspects of problem solving. Sometimes the client does not have the precise data for the problem at hand, and provides only some imprecise estimations. However,

*Corresponding author

imprecise input can only lead to imprecise output, and the resulting solution might not be optimal. Different approaches have been proposed to deal with this situation, such as robust optimization or stochastic optimization, see [1, 2, 3] for surveys on those areas.

In this paper we follow the paradigm of *optimizing under explorable uncertainty*, which is an alternative approach which has been introduced in 1991 [4], and started to be applied to scheduling problems in 2016 [5]. In this approach, a problem instance consists of a set of numerical parameters, the algorithm obtains as input only an uncertainty interval for each one. The algorithm knows for each parameter that it belongs to the given interval and has the possibility to make a query in order to obtain the precise value. Clearly a compromise has to be found between the number of queries an algorithm makes and the quality of the solution it produces. This setting differs from a probabilistic one, studied in [5, 6], where jobs have weights and processing times drawn from known distributions, and the algorithm can query these parameters. A seemingly similar problem has been studied in [7, 8, 9], where the term *testing* has a different meaning than in this paper.

Different measures to evaluate the performance of an algorithm under explorable uncertainty have been investigated. Early work studied the number of queries required in order to be able to produce an optimal solution, no matter what the values of the non queried parameters happen to be. In this sense the queries are used to form a proof of optimality, and the underlying techniques are close to the ones used in query complexity.

A broad range of problems has been studied, namely the finding the median [10], or more generally the k -th smallest value among the given parameters [11], sorting [12], finding a shortest path when the uncertain parameters are the edge lengths [13], determining the convex hull of given uncertain points [14], or computing a minimum spanning tree [15, 16, 17]. The techniques which have been developed in these papers have been generalized and described in [18, 19].

A related stochastic model for scheduling with testing has been introduced by Levi, Magnanti and Shaposhnik [6, 20]. They consider the problem of minimizing the weighted sum of completion times on one machine for jobs whose processing times and weights are random variables with a joint distribution, and are independent and identically distributed across jobs. In their model, testing a job provides information to the scheduler (by revealing the exact weight and processing time for a job, whereas initially only the distribution is known). They present structural results about optimal policies and efficient optimal or near-optimal solutions based on dynamic programming.

2. Our Contribution

Before defining formally the problem, we start with an illustrative example. Suppose that you need to schedule on a single machine four jobs called A, B, C, D . Job B has processing time 5, while jobs A, C, D have processing times 0.3, see Figure 1. The cost of a schedule (also called *objective value*) is the total completion time of the jobs. Hence one possible optimal execution




Figure 1: Some schedules with four jobs A, B, C, D and parameters $p = 0.3, x = 4.7$. White boxes represent job tests, while gray boxes represent job executions. The cost is the sum of completion times or equivalently the sum of the execution lengths of the boxes, multiplied with their respective ranks.

order is A, C, D, B , which has objective value 7.7. However the processing times are initially unknown to you, the jobs are indistinguishable, and you only know that the processing times are either 0.3 or 5. Hence you have no information to distinguish among different execution orders. For example if you schedule them in order of their identifiers, then in our example the execution order A, B, C, D would yield an objective value of 17.1, which is roughly 2.22 larger than the optimum. Luckily you have access to a processing time oracle, which allows you to query the processing time of a particular job, and this operation takes 1 time unit. Such an oracle can be thought as a machine learning black box predictor, which was trained on large number of jobs. Throughout the paper we use indistinguishable the word *test* or *query* for this operation. You could query this oracle for all 4 jobs, providing full information, which allows you to schedule the

jobs in the optimal order. However you don't have to wait until you know the processing time of all jobs before you start executing them. For example if a query reveals a short job, then it could be scheduled immediately, and if a query reveals a long job, then its execution could be postponed towards the end of the schedule. This means that there is no benefit to query the last job, as the position of its execution would be the same, regardless of the outcome. In summary, for this example you have the possibility to query between 0 and 3 jobs. The resulting schedules are depicted in Figure 1. Testing the first 2 jobs would lead to the smallest objective value of 14.7, which is only roughly 1.91 larger than the optimum.

The general problem studied in this paper consists in scheduling n independent jobs on a single machine, with jobs duration limited to two possible values, p and $p+x$, for some given parameters $p, x > 0$. Each job can be tested, requiring one time unit on the schedule, revealing the actual processing time. The objective value of a schedule is the sum of the completion times of the jobs, and is also called the *cost* of the schedule. This value is normalized by dividing it with the minimum cost over all schedules, which is called the *optimum*. Computing the optimum requires full knowledge of the job's processing times. This ratio is called the *competitive ratio*, and the goal is to design an algorithm with smallest possible competitive ratio. The ratio measures the price of not knowing initially all processing times.

We emphasize on the importance of normalizing the objective value. The worst case instance for the sum of completion times, consists of only long jobs, and any test would strictly increase the objective value. In other words, if the goal were to minimize the objective value in the worst case (and not the competitive ratio), then the best strategy would execute all jobs untested.

Competitive analysis has been introduced during the 1980's to study computational models where the input is a request sequence given sequentially to the algorithm [21]. Our problem does not fall exactly into this setting, nevertheless we borrow its terminology.

Competitive ratio is usually seen as the value of a game played between an algorithm deciding which jobs to test and an adversary deciding the job lengths, hence generating the instance. This is a zero sum game, and therefore it admits an optimal algorithm and a schedule produced at equilibrium.

We study the competitive ratio in two algorithmic models. In the *non-adaptive model*, the algorithm has to decide at the beginning of the schedule how many jobs it wants to test, while in the *adaptive model*, it can make this decision adaptively, depending on the outcome of previous tests.

At the beginning of the schedule, jobs are indistinguishable to the algorithm, therefore we suppose that the algorithm processes the jobs in a fixed arbitrary order. As we show in Lemma 2, an optimal algorithm executes a tested short job right after its test, and delays the execution of a tested long job towards the end of the schedule.

With this observation in mind, the behavior of an algorithm can be fully described by a binary decision for each job, namely to execute it untested, or to test it and to execute it accordingly to the outcome as described above. From experiments, it seems that optimal algorithms follow a *two-phase strategy*, namely to first test some number of jobs, and then to execute the remaining jobs untested (see Section 8). This does not specify where the tested jobs are executed, even though the optimal algorithm executes them in a specific manner (see previous paragraph and Lemma 2).

In the example of Figure 1 the first two jobs are tested and the the last two jobs are executed untested. Intuitively, this makes sense, as the benefit of a test decreases with the progression of the schedule. The purpose of a test is to identify long jobs so they can be scheduled as late as possible. This allows to decrease the objective value, but the benefit depends on the number of jobs whose processing time is not yet known. In Section 5 we were able to provide a formal proof for the last 3 jobs in the schedule, which however does not generalize. Hence we leave the proof of this conjectured dominant behavior as an open problem, and focus in this paper only on two-phase strategies.

Conjecture 1. *For all values of p, x, n , and for both the adaptive and non-adaptive models, there is an optimal algorithm following a two-phase strategy.*

Here by *optimal algorithm*, we mean an algorithm which achieves the smallest competitive ratio for the worst case instance, i.e. when the adversary plays optimally. In case the adversary does not play optimally, in the adaptive model the algorithm can achieve an even better competitive ratio, but might need to diverge from a two-phase strategy for this.

For the non-adaptive model, we were able to provide an algorithm running in time $O(n^2)$ which determines the optimal two-phase strategy. Note that this procedure is not polynomial in the input size, which consists only of the values p, x, n , but it is polynomial in the number of jobs, and therefore also in the size of the produced schedule. In addition we provide a closed form expression of the limit of the competitive ratio, when n tends to infinity.

For the adaptive model, we were also able to provide an algorithm running in time $O(n^3)$. More precisely, if there are still r jobs to be handled, the algorithm computes in time $O(r^3)$ a strategy for the remainder (how many jobs it wants to test). The algorithm will stick to this strategy if the adversary behaves optimally. However once the adversary diverges from its optimal strategy, the algorithm needs to recompute the strategy for the remaining jobs. This means that if the algorithm plays against the optimal adversary it would spend time $O(n^3)$ to decide whether to test the first job, and $O(1)$ for every subsequent job, otherwise it has time complexity $O(n^3)$ on every job decision.

3. Problem settings

Formally, we consider the problem of scheduling n independent jobs on a single processor with the objective of minimizing the sum of completion times. Jobs are numbered from 1 to n . Every job can be either short (processing time p) or long (processing time $p + x$) for some known parameters $p, x > 0$. The algorithm receives n jobs without the information of their processing times and will handle them in order of their indices. For every job, the algorithm has the choice to test it or to execute it untested. A job test consists in the execution of a processing time oracle whose duration is one time unit and reveals to the algorithm the processing time of the job. In principle a tested job could be scheduled at any moment in the schedule, but as we show in Lemma 2, it is dominant to schedule short jobs immediately after their test and to postpone long tested jobs towards the end of the schedule.

An example is pictured in Figure 1. Here, the scheduler decides to test the first two jobs. The first test indicates a short job that is executed immediately, and the second test indicates a long job that is postponed to the end of the schedule. The last two jobs are executed untested.

Borrowing the terminology of online algorithms we consider the problem as a game played between an algorithm and an adversary. See [22, 23] for recent introductions into online algorithms. But we emphasize that this is not an online problem, in the sense that the algorithm receives from the beginning on all n jobs. It chooses a strategy $u \in \{T, E\}^n$, while the adversary chooses a strategy $v \in \{p, x\}^n$. Here we abuse notation and use symbols T, E, p, x , even though p, x are also used for the numerical parameters of the game. For every job $j \in \{1, \dots, n\}$, if $u_j = T$, then the algorithm tests the job, otherwise the algorithm executes the job untested. If $v_j = p$ then job j is short, otherwise it is long. The resulting competitive ratio is a function of the strategies u and v , and is formally defined by the pseudocode given in Algorithm 1. A two-phase strategy u has the shape T^*E^* , using a regular expression notation.

The algorithm wants to minimize the competitive ratio, while the adversary wants to maximize it. In the non-adaptive model, the algorithm plays u , and the adversary plays v in response, while in the adaptive model, algorithm and adversary play in alternation.

For convenience we describe a schedule with the string $u_1v_1 \dots u_nv_n$, which encodes the decisions made by both the algorithm and the adversary, and compactly describes the actual schedule. Such a string matches the regular expression $((T|E)(p|x))^n$. For example the schedule resulting of testing 2 jobs in Figure 1 would be described by $TpTxEpEp$. When both the algorithm and the adversary play optimally, then we call the resulting schedule the *equilibrium schedule*. In this notation Ep describes the execution of an untested short job, Ex the execution of a long untested job, Tp the test of a short job, followed immediately by its execution and Tx the test of a long job, whose execution is delayed towards the end of the schedule and does not explicitly appear in the notation.

Algorithm 1 Pseudocode defining the competitive ratio as a function of the strategies $u \in \{T, E\}^n, v \in \{p, x\}^n$. ALG represents the cost of the schedule produced by the algorithm and OPT the optimal cost.

```

ALG =  $pn(n+1)/2$  ▷ base cost
 $r = n$  ▷ rank of next job to be processed
for  $j = 1, \dots, n$  do ▷ Process jobs in given order
  if  $u_j = E$  and  $v_j = p$  then ▷ execute short untested job
     $r = r - 1$ 
  end if
  if  $u_j = E$  and  $v_j = x$  then ▷ execute long untested job
    ALG = ALG +  $rx$ 
     $r = r - 1$ 
  end if
  if  $u_j = T$  and  $v_j = p$  then ▷ test short job and execute
    ALG = ALG +  $r$ 
     $r = r - 1$ 
  end if
  if  $u_j = T$  and  $v_j = x$  then ▷ test long job and postpone
    ALG = ALG +  $r$ 
  end if
end for
ALG = ALG +  $xr(r+1)/2$  ▷ execute postponed long tested jobs
 $\ell = |\{j : v_j = x\}|$  ▷ number of long jobs
OPT =  $(pn(n+1) + x\ell(\ell+1))/2$ 
return ALG / OPT

```

Rank

In this paper it will often be convenient to express the cost of a schedule — the total completion time — using the notion of *rank*. A schedule consists of a sequence of job tests and job executions. Each of these actions has a rank which is defined as the number of jobs which are executed after this part, including the job itself in case of a job execution. This permits us to express the total job completion time as the sum of all actions in the schedule of the length of the action multiplied by its rank. For example in Figure 1 (algorithm testing two jobs), the first test duration delays the completion time of all 4 jobs, it thus has rank 4. The second test delays all but the first job, therefore it has rank 3. The tested long job is postponed to the end of the schedule, and has rank 1. The two last jobs which are executed untested have rank respectively 3 and 2.

Two algorithmic models

We study two models. In the *adaptive model*, the algorithm can adapt to the adversary after each step. This setting is studied in Section 7. In contrast, in the *non-adaptive model* the algorithm has to decide once for all on a sequence of testing and executing, and stick to it, no matter what the job lengths happen to be. The results in this setting are presented in Section 6.

See Table 1 for an illustration of the non-adaptive model. The algorithm chooses a particular strategy (column) and the adversary chooses a particular response (row). The resulting ratio is indicated in the selected cell of this array. The equilibrium schedule is determined by the min-max value of this array, namely *Exp* in this case, resulting in the ratio $11/7$.

In the adaptive model, the interaction between the algorithm and the adversary is illustrated by the game tree shown in Figure 2. Every node represents a particular moment of the interaction. Leaf nodes are labeled with the resulting ratio. Inner nodes have two out-going arcs, representing the possible actions of the algorithm or the adversary, which play in alternation, and are labeled with the ratio resulting of a best choice. The algorithm is the minimizer in this game, while the adversary is the maximizer. The root is labeled with the competitive ratio of the game, which is $11/7$ in our example. For this small example, the equilibrium schedule happen to be the same in both the adaptive and the non-adaptive model, but for larger number of jobs, the game tree would be too big to be readable.

4. Dominance properties

The algorithm processes jobs in given order. At any moment it can decide to test the next job in this order, or to execute it untested, or to execute a previously tested job. In this section we show some dominant behavior of optimal algorithms.

We say that an algorithm is *postponing* if it executes all tested long jobs at the end of the schedule and executes every tested short job immediately after its test. We will show that any

	EE	ET	TE	TT
pp	$\frac{3p}{3p} = 1$	$\frac{3p+1}{3p} = \frac{4}{3}$	$\frac{3p+2}{3p} = \frac{5}{3}$	$\frac{3p+3}{3p} = 2$
px	$\frac{3p+x}{3p+x} = 1$	$\frac{3p+x+1}{3p+x} = \frac{8}{7}$	$\frac{3p+x+2}{3p+x} = \frac{9}{7}$	$\frac{3p+x+3}{3p+x} = \frac{10}{7}$
xp	$\frac{3p+2x}{3p+x} = \frac{11}{7}$	$\frac{3p+2x+1}{3p+x} = \frac{12}{7}$	$\frac{3p+x+2}{3p+x} = \frac{9}{7}$	$\frac{3p+x+4}{3p+x} = \frac{11}{7}$
xx	$\frac{3p+3x}{3p+3x} = 1$	$\frac{3p+3x+1}{3p+3x} = \frac{16}{15}$	$\frac{3p+3x+2}{3p+3x} = \frac{17}{15}$	$\frac{3p+3x+4}{3p+3x} = \frac{19}{15}$

Table 1: The competitive ratio in the non-adaptive setting for 2 jobs, $p = 1$ and $x = 4$, as a function of the algorithm's strategy (columns) and the adversarial strategy (rows).

optimal algorithm has this property. For the proof we need to show that once a tested job turns out to be long, then it is better to postpone it rather to execute it at an earlier time in the schedule. It might be possible to compare the costs of the algorithm using an exchange argument, which is a standard technique in scheduling. The problem however, is that the adversary can behave differently in both situations, and we need to compare the ratio not the actual cost of the algorithm. Hence the proof is more complicated and involves the use of a restricted adversary, allowing us to reduce to comparable situations.

We say that the adversary is *restricted* if it is forced to make the last job long. The optimal adversary does not behave that way, but this restriction is needed for the proof. Any schedule ends with a (possibly empty) part consisting of the execution of some tested long jobs. Those jobs are said to be *postponed*. We consider the game tree resulting of an adaptive algorithm (not necessarily postponing) playing against an optimal adversary. For every node where the algorithm is in turn, 3 actions are possible, namely to test a job, to execute an untested job or to execute a tested long job. After the first two actions the adversary is in turn, while after the last action the algorithm is still in turn. Figure 2 shows only the first two actions.

Every node has a ratio associated to it, which is the competitive ratio reached from this node, when both the algorithm and the adversary play optimally. Consider a node in the game tree, where the algorithm is in turn, and assume that the algorithm did not yet execute untested jobs. The ratio of this node is completely described by some parameters c, d, e, f : c is the number of tested short jobs, d the number of tested long jobs which are pending (not yet executed), f is the number of untested jobs, and e is a value determining the algorithm's cost in the following way. It is defined as the total rank over all tests plus x times the total rank of all executions of long

tested jobs which are not postponed. If this node is a leaf in the game tree, the algorithm's cost is $e + xd(d+1)/2 + pn(n+1)/2$, where we added to e the additional cost of the postponed long job executions, and a *base cost* generated by all n job executions. We use the following notations for the ratio of a node with parameters c, d, e, f , under various restrictions on the players.

notation	algorithm	adversary
$\text{ratio}(c, d, e, f)$	unrestricted	unrestricted
$\text{ratio}'(c, d, e, f)$	postponing	unrestricted
$\text{ratio}''(c, d, e, f)$	postponing	restricted

We can show the following useful inequality. It says that if the algorithm is artificially provided with the knowledge that one of the untested jobs is long, then this information provides an advantage to the algorithm which drops the competitive ratio.

Lemma 1. *For any $c, d \geq 0$, $e > 0$ and $f \geq 1$ we have $\text{ratio}''(c, d, e, f) > \text{ratio}'(c, d+1, e, f-1)$.*

Proof. First we observe $\text{ratio}''(c, d, e, f) \leq \text{ratio}'(c, d, e, f)$, which holds simply because a restricted adversary has less choices than an unrestricted adversary. More arguments are necessary to show that the inequality is strict. Hence consider the moment t when the algorithm processes the last job (either by testing it or by executing it untested). No matter what the action of the algorithm is for this job, the job will be executed before all d pending long jobs. Then if the adversary makes this last job short instead of long, the cost of the algorithm decreases by exactly xd , while the cost of the adversary decreases by at least xd , in fact by $x(d+d')$, where d' is the number of long jobs which have been executed by the algorithm before t . As a result, the competitive ratio is strictly increasing, unless the algorithm and the adversary produce exactly the same schedules. This would mean that the algorithm tested no job, and executed no long job before time t , which contradicts the assumption $e > 0$.

Next we observe the equality $\text{ratio}''(c, d, e, f) = \text{ratio}'(c, d+1, e, f-1)$. The point is that the last processed job j executed in the equilibrium schedule corresponding to $\text{ratio}''(c, d, e, f)$ is long and followed immediately by the d tested postponed long jobs. Hence conceptually one could consider this job as an additional tested long job, decreasing at the same time the number of untested jobs. Because the restricted adversary was only committed to job j , and not to the other $f-1$ jobs, the equilibrium schedule corresponding to $\text{ratio}''(c, d, e, f)$ is identical to the one corresponding to $\text{ratio}'(c, d+1, e, f-1)$. This concludes the proof. \square

The previous lemma allows us to show that the optimal algorithm is postponing.

Lemma 2. *Without loss of generality the optimal algorithm is postponing, in other words $\text{ratio} = \text{ratio}'$. This holds for both the adaptive and the non-adaptive model and under Conjecture 1.*

Proof. We provide the proof only for the adaptive model, as it is the hardest. First we observe that a tested short job j can be safely executed by the algorithm right after its test. This choice reduces the rank of the following actions happening between this moment and the eventual execution of job j , and has no influence on the adversary's strategies. Hence the optimal algorithm executes tested short jobs right after their test.

For the tested long jobs, formally we show that for every node in the game tree with parameters c, d, e, f , the optimal algorithm will not execute a tested long job as the next action. The proof is by induction on the pair (f, d) in lexicographical order. The base case $f = 0$, holds trivially. Indeed the node is in fact a leaf in the game tree, and the algorithm has no option but to execute all d tested long jobs.

For the induction step consider a node with parameters c, d, e, f such that $f \geq 1$ and assume the induction hypothesis for all nodes with parameter $f - 1$, or with parameter f and $d - 1$ in case $d \geq 1$.

We start with the easy case $d = 0$. When there are no tested long jobs available, the algorithm has no choice but to test or to execute the next one among the f untested jobs. This action results in a node with parameter $f - 1$, and by induction hypothesis the algorithm is postponing.

It remains to show the induction step for the case $d \geq 1$. The following table shows the ratios resulting by an action from the algorithm followed possibly by an action from the adversary.

short	action	ratio
Tp	test a short job	$\text{ratio}(c + 1, d, e + f + d, f - 1)$
Tx	test a long job	$\text{ratio}(c, d + 1, e + f + d, f - 1)$
Ep	execute an untested short job	$C = \text{ratio}(c + 1, d, e, f - 1)$
Ex	execute an untested long job	$B = \text{ratio}(c, d, e + x(f + d), f - 1)$
	execute a tested long job	$A = \text{ratio}(c, d - 1, e + x(f + d), f)$

Ratios A, B can be compared using the fact that by induction hypothesis for nodes with parameters $f - 1, d$ or $f, d - 1$ the optimal algorithm is postponing, which allows to use Lemma 1, with $d - 1$ instead of d to match the expressions. It shows that ratio A is larger than ratio B .

But this is not enough, we need to show $A \geq \max\{B, C\}$. This inequality would certify that the optimal algorithm does not choose to execute a tested long job now, which is the induction step to show.

When $B \leq C$ there is nothing to show. Hence we assume $C > B$, in other words the optimal adversary answers the E action (execute an untested job) by a short job. We claim that from now on all executions of untested jobs will be short. Here we use the two-phase assumption. This assumption states that once an untested job is executed, all remaining untested jobs will be executed untested as well. As a result the subsequent schedule consists of f job executions

followed by d executions of the pending tested long jobs. The two-phase assumption really means that from now on the algorithm commits to this execution pattern, hence the adversary can decide on a number $0 \leq b \leq f$ of jobs to be long. Which of the f executed jobs will be long is of no influence on the cost of the adversary, but placing them first maximizes the cost of the algorithm. This proves the claim.

Now if the algorithm knows that its remaining f executions of untested jobs will be short, then we claim that it prefers to execute an untested job, rather than a tested long job. The formal reason is that if the algorithm decides for the later, then it will have a ratio strictly larger than C already if the adversary still decides to make all untested jobs short, no matter if they are tested or executed untested.

This concludes the induction step, and therefore the proof. \square

We complete this section by describing what is happening at the last node in the game tree.

Lemma 3. *Consider a node in the game tree where the algorithm is in turn and a single untested job j is left, that is a node with parameters c, d, e and $f = 1$. Then the optimal algorithm executes job j untested, following by all d tested long jobs. Moreover the adversary will make job j short. This holds for both the adaptive and the non-adaptive model and does not rely on Conjecture 1.*

Proof. If the adversary makes j short instead of long, then it decreases the cost of the algorithm by $x(d+1)$ at most (depending on the rank of j) and decreases his cost by at least $x(d+1)$. Hence making j short increases strictly the competitive ratio and is therefore the choice of the optimal adversary. Here it is important that this was the last interaction between the algorithm and the adversary.

But since j will be short anyway, the algorithm does not get any benefit from testing it. Executing j untested saves the delay caused by the test. Moreover the optimal execution order for the algorithm is to start executing job j followed by the d pending long jobs. \square

5. Argument in favor of Conjecture 1

In this section we study the last actions conducted by the algorithm in an equilibrium schedule. By the previous lemma, the last action is E , i.e. executing the last job untested. We show that if the second last action is T , then the third last action must be T as well. A proof for Conjecture 1 would consist of a similar statement for arbitrary positions in the schedule. However in the proof below, it is crucial that there is no more interaction between the algorithm and the adversary after this position. Otherwise we would need precise bounds on the costs of the remaining schedule. This in turn would mean that we would be able to provide precise bounds on the competitive ratio as a function of n, p and x . But those bounds are extremely difficult to obtain in the adaptive model.

Lemma 4. *Consider a node in the game tree where the algorithm is in turn and 3 untested jobs are left, that is a node with parameters c, d, e and $f = 3$. If the optimal algorithm chooses strategy E at this node, then he will choose strategy E in the next node as well.*

Proof. By Lemma 3, the remaining interaction between the algorithm and the adversary concerns the third and second last job only. There are 16 possible combined strategies in this interaction, which we enumerate in Table 2. For the strategy $EpEpEp$, we denote ALG_0 the cost of the algorithm and OPT_0 the cost of the adversary. For the 15 remaining strategies, we formulate the respective costs by adding terms to these base costs. Table 2 summarizes these costs. For convenience we denote by *node* w the node in the game tree reached by following the strategy w , which can be any prefix of one of the 16 strategies listed in Table 2. In addition, throughout the proof, we denote by $\text{ratio}(w)$ the competitive ratio obtained when the strategy is w , where w is any of the 16 considered strategies.

strategy	$ALG_0 + \dots$	$OPT_0 + \dots$
$EpEpEp$	0	0
$EpExEp$	$x(d+2)$	$x(b+d+1)$
$EpTpEp$	$+d+2$	0
$EpTxEp$	$x(d+1) + d + 2$	$x(b+d+1)$
$ExEpEp$	$x(d+3)$	$x(b+d+1)$
$ExExEp$	$x(2d+5)$	$x(2b+2d+3)$
$ExTpEp$	$x(d+3) + d + 2$	$x(b+d+1)$
$ExTxEp$	$x(2d+4) + d + 2$	$x(2b+2d+3)$
$TpEpEp$	$+d+3$	0
$TpExEp$	$x(d+2) + d + 3$	$x(b+d+1)$
$TpTpEp$	$+2d+5$	0
$TpTxEp$	$x(d+1) + 2d + 5$	$x(b+d+1)$
$TxEpEp$	$x(d+1) + d + 3$	$x(b+d+1)$
$TxExEp$	$x(2d+4) + d + 3$	$x(2b+2d+3)$
$TxTpEp$	$x(d+1) + 2d + 6$	$x(b+d+1)$
$TxTxEp$	$x(2d+3) + 2d + 6$	$x(2b+2d+3)$

Table 2: Costs of the algorithm and the optimal cost for the 16 considered strategies.

In the remainder of the proof, we study optimal strategies for either player in different situations.

Answer on ExT is p

We consider the case when the algorithm executes a job, which the adversary makes long, and then tests a job. We will show that the adversary will make the tested job short. For a proof by contradiction assume

$$\begin{aligned} & \text{ratio(ExTpEp)} < \text{ratio(ExTxEp)} \\ \Leftrightarrow & \frac{\text{ALG} + x(d+3) + d + 2}{\text{OPT} + x(b+d+1)} < \frac{\text{ALG} + x(2d+4) + d + 2}{\text{OPT} + x(2b+2d+3)} \\ \Leftrightarrow & \frac{\text{ALG} + x(d+3) + d + 2}{\text{OPT} + x(b+d+1)} < \frac{x(d+1)}{x(b+d+2)}. \end{aligned}$$

But left hand side is a competitive ratio and hence at least 1, while right hand side is less than 1. We reached a contradiction. Note that for the last inequality we used the implication $\alpha/\delta < (\alpha + \beta)/(\delta + \gamma) \Leftrightarrow \alpha/\delta < \beta/\gamma$ which holds for any $\delta, \gamma > 0$.

Answer on EpT is p

Similarly, for a proof by contradiction assume

$$\begin{aligned} & \text{ratio(EpTpEp)} < \text{ratio(EpTxEp)} \\ \Leftrightarrow & \frac{\text{ALG} + d + 2}{\text{OPT}} < \frac{\text{ALG} + x(d+1) + d + 2}{\text{OPT} + x(b+d+1)} \\ \Leftrightarrow & \frac{\text{ALG} + d + 2}{\text{OPT}} < \frac{x(d+1)}{x(b+d+1)}. \end{aligned}$$

But again, left hand side is at least 1, while right hand side is at most 1. We reached a contradiction.

Answer on Ex is E

Suppose that the answer on *ExE* is *p*. Then at node *Ex* the algorithm needs to decide between *ExTpEp* and *ExEpEp*. But the later has smaller ratio, and will be chosen.

So assume from now on that the answer on *ExE* is *x*. To show that the answer on *Ex* is *E*, we assume the contrary, namely that the answer is *T*. In other words we assume the following inequality.

$$\begin{aligned} & \text{ratio(ExTpEp)} < \text{ratio(ExExEp)} \\ \Leftrightarrow & \frac{\text{ALG} + x(d+3) + d + 2}{\text{OPT} + x(b+d+1)} < \frac{\text{ALG} + x(2d+5)}{\text{OPT} + x(2b+2d+3)} \\ \Leftrightarrow & \frac{\text{ALG} + x(d+3) + d + 2}{\text{OPT} + x(b+d+1)} < \frac{x(d+2) - d - 2}{x(b+d+2)}. \end{aligned}$$

As in the previous cases, the left hand side is at least 1, and the right hand side less than 1, leading to a contradiction.

Answer on ExE is p

For a proof by contradiction assume the adversary answers x on ExE (breaking ties towards p). This translates into the following inequality.

$$\begin{aligned} \text{ratio}(\text{ExExEp}) &> \text{ratio}(\text{ExEpEp}) \\ \Leftrightarrow \frac{\text{ALG} + x(2d + 5)}{\text{OPT} + x(2b + 2d + 3)} &> \frac{\text{ALG} + x(d + 3)}{\text{OPT} + x(b + d + 1)} \\ \Leftrightarrow \frac{d + 2}{b + d + 2} &> \frac{\text{ALG} + x(d + 3)}{\text{OPT} + x(b + d + 1)}. \end{aligned}$$

We reach the same contradiction, as in all previous cases.

Answer on Ep is E

Suppose that the answer on EpE is p. Then at the node Ep the algorithm needs to decide between EpTpEp and EpEpEp. But the later has smaller ratio, and will be chosen.

So assume from now on that the answer on EpE is x. Hence we have just 3 possible strategies to consider, namely ExEpEp, EpTpEp and EpExEp. If the adversary answers the first E by x, then the ratio will be ratio(ExEpEp), while if the adversary answers by p, then the ratio will be at most ratio(EpExEp).

However ExEpEp has larger ratio than EpExEp, as the algorithm has larger cost while the adversary has the same cost. We conclude that the adversary answers the first E by p.

All cases have been covered in this analysis, which concludes the proof of the lemma. \square

6. Non adaptive algorithms

In this section we analyze algorithms in the non-adaptive model following a two-phase strategy. The set of algorithmic strategies can be modeled as $u \in \{T^a E^{n-a} : 0 \leq a \leq n\}$, and the set of adversarial strategies modeled as $v \in \{x, p\}^n$. The resulting schedule is described by the string $u_1 v_1 u_2 v_2 \dots u_n v_n$.

In addition we show some structure of the adversarial best response to any algorithmic strategy.

Lemma 5. *The adversarial best response v to any algorithmic strategy u of the form¹ $T^* E^*$ results in a schedule of the form $(Tx)^*(Tp)^*(Ex)^*(Ep)^*$.*

Proof. We show the claim by an exchange argument. Let v be an adversarial strategy which is not of the claimed form. Then there is a position $0 \leq i < n$ such that $u_i v_i u_{i+1} v_{i+1} \in \{TpTx, EpEx\}$. We consider the effect of swapping v_i with v_{i+1} in the adversarial strategy. Since the number of long jobs is preserved, in order to analyze the change in the ratio, in fact we need to analyze the

¹We use regular expressions to describe the form of a strategy and the form of a schedule.

cost of the resulting schedule. We observe that the cost is increased by the exchange, namely by x in the case $EpEx$ and by 1 in the case $TpTx$. \square

The implication of this lemma is that there are only $O(n^3)$ possible outcomes of the game. One can compute the ratios of these schedules in amortized constant time by considering them in an appropriate order and updating the costs of the algorithm's schedule and optimal schedule during the loops. Hence computing the optimal non-adaptive strategy of the form T^*E^* can be done in time $O(n^3)$. We show how this complexity can be reduced.

algorithm:	Tx	Tp	Ex	Ep	x
part	1	2	3	4	5
number	d	$a - d$	$\ell - d$	$n - a + d - \ell$	d
fraction	β	$\alpha - \beta$	γ	$1 - \alpha - \gamma$	β

optimum:	Ep	Ex
fraction	$1 - \beta - \gamma$	$\beta + \gamma$

Figure 3: Schedule obtained in the non-adaptive setting.

Theorem 1. *For the non-adaptive model, the equilibrium schedule (and therefore the optimal algorithm's strategy) can be computed in time $O(n^2)$.*

Proof. We observed earlier that all equilibrium schedules consist of four parts and are of the form $(Tx)^d(Tp)^{a-d}(Ex)^{\ell-d}(Ep)^{n-a+d-\ell}$ for some parameters a, d, ℓ , followed by a fifth part executing the postponed long jobs, see Figure 3. The parameter a describes the number of tests done by the algorithm, among which d describes the total number of long tested jobs. In addition ℓ denotes the total number of long jobs. We fix the parameters a, ℓ and show that the optimal parameter d for the adversary can be computed in constant time. As there are only $O(n^2)$ possible values for a and ℓ , this would prove the theorem.

The parameter ℓ defines the total number of long jobs, hence the optimal schedule, which is $(Ep)^{n-\ell}(Ex)^\ell$ and does not depend on d . As a result, the adversary chooses d such that the cost

of the algorithm is maximum. The cost of the algorithm can be expressed as follows

$$\begin{aligned}
\text{ALG}(d) = & d(d+1)/2 && \text{(cost of part 1)} \\
& + dn && \text{(delay caused by part 1)} \\
& + (1+p)(a-d)(a-d+1)/2 && \text{(cost of part 2)} \\
& + (1+p)(a-d)(n-a+d) && \text{(delay caused by part 2)} \\
& + (p+x)(\ell-d)(\ell-d+1)/2 && \text{(cost of part 3)} \\
& + (p+x)(\ell-d)(n-a+2d-\ell) && \text{(delay caused by part 3)} \\
& + p(n-a+d-\ell)(n-a+d-\ell+1)/2 && \text{(cost of part 4)} \\
& + p(n-a-\ell+d)d && \text{(delay caused by part 4)} \\
& + (p+x)d(d+1)/2. && \text{(cost of delayed long jobs)}
\end{aligned}$$

This function is quadratic in d with a negative second derivative, namely $-2x$. Therefore the integer maximizer of ALG can be found, by first finding the root of the function $d \mapsto \text{ALG}(d) - \text{ALG}(d-1)$ and then rounding it down. This function is $a + (a - 2d + 2\ell - n + 1)x$, and its root is

$$d^* = \ell - \frac{n - a - a/x - 1}{2},$$

which never exceeds ℓ . The adversary has to choose d , such that $0 \leq d \leq \min\{a, \ell\}$, hence he chooses

$$\min\{a, \max\{0, \lfloor d^* \rfloor\}\}.$$

Since this integer can be computed in constant time, this concludes the proof. The optimal algorithm is summarized in Figure 2. \square

We conclude this section, by providing a closed expression for the asymptotic competitive ratio. The computations have been conducted with the help of the computing system Mathematica. The files are available to the reader in the companion webpage <https://www.lip6.fr/Christoph.Durr/Scheduling-w-Oracle/>.

Formally we analyze the limit of the competitive ratio when n tends to infinity. This is possible, because for large values of n , we can focus on the dominant parts in the schedule costs, and ignore integrality of the parameters. In the sequel instead of working with integral parameters a, d, ℓ we work with fractions. Note that along the way we also made a variable change, in the sense $a = \alpha n, d = \beta n, \ell = (\beta + \gamma)n$, see Figure 3.

The intuition behind the following theorem is that for large instances, if x is not larger than $2 + 1/p$ then the difference of short and large jobs is so small that the length of a test is too costly compared to the gain obtained from the returned information.

Algorithm 2 Pseudocode computing the optimal strategy for the algorithm in the non-adaptive model.

```

best = +∞
a* = None
for all 0 ≤ a, ℓ ≤ n do
    OPT = pn(n + 1)/2 + xℓ(ℓ + 1)/2
    d* = ℓ - (n - a - a/x - 1)/2
    d = min{a, max{0, ⌊d*⌋}}
    Ratio = ALG(d)/OPT
    if Ratio < BestRatio then                                ▷ keep best ratio
        BestRatio = Ratio
        a* = a
    end if
end for
return “The optimal strategy tests the first a* jobs.”

```

Theorem 2. *The asymptotic competitive ratio of the scheduling problem for non-adaptive algorithms is*

$$\sqrt{1 + \frac{x}{p}}$$

when $0 \leq x < 2 + 1/p$ and

$$1 + \frac{x^2 - px - 1 + \sqrt{\Delta'}}{2px^2}$$

when $x \geq 2 + 1/p$, for $\Delta' = 8p(x - 1)x^2 + (1 + px - x^2)^2$.

Proof. Given n , the number of jobs, the algorithm decides on some fraction $0 \leq \alpha \leq 1$, and will test the first αn jobs, followed by the untested execution of the remaining $(1 - \alpha)n$ jobs. The adversary decides on some fraction $0 \leq \beta \leq \alpha$ of the tested jobs to be long, and some fraction $0 \leq \gamma \leq 1 - \alpha$ of the executed jobs to be long. Both the cost of the algorithm and of the optimal schedule are quadratic expressions in n . For the sake of simplicity we focus only on the quadratic part, which is dominating for large n . Hence the results of this section apply only asymptotically when n tends to infinity. However in principle it is possible to make a similar but tedious analysis also for the linear part and obtain results that hold for all job numbers n .

The game played between the algorithm and the adversary contains the following strategies. The algorithm chooses some parameter α while the adversary chooses some parameters β, γ . The

n^2 dependent part of the cost of the algorithm's schedule is (we multiply by 2 to simplify notation)

$$2 \cdot \text{ALG} = 2\beta \tag{1}$$

$$+(1+p)(\alpha-\beta)^2 \tag{2}$$

$$+2(1+p)(\alpha-\beta)(1-\alpha+\beta) \tag{3}$$

$$+(p+x)\gamma^2 \tag{4}$$

$$+2(p+x)\gamma(1-\alpha+\beta-\gamma) \tag{5}$$

$$+p(1-\alpha-\gamma)^2 \tag{6}$$

$$+2p(1-\alpha-\gamma)\beta \tag{7}$$

$$+(p+x)\beta^2. \tag{8}$$

We justify the cost, by distinguishing the contribution of each of the parts of the schedule to its cost, referring to the 5 parts of the schedule illustrated in Figure 3. Expression (1) is the delay caused by part 1 on the following parts of the schedule, (2) is the cost of part 2, (3) is the delay caused by part 2 on the rest, (4) is the cost of part 3, (5) is the delay of part 3 on the rest, (6) is the cost of part 4, (7) is the delay of part 4 on part 5, (8) is the cost of part 5.

The n^2 dependent part of the optimal schedule (again multiplying by 2 to simplify notations) is

$$2 \cdot \text{OPT} = p + x(\beta + \gamma)^2.$$

The algorithm wants to minimize the ratio ALG/OPT while the adversary wants to maximize it. In order to avoid the manipulation of fractions, we introduce another parameter r chosen by the algorithm. It defines the expression

$$\begin{aligned} G &= (1+r)2 \cdot \text{OPT} - 2 \cdot \text{ALG} \\ &= pr + \alpha^2 + \beta^2 - 2\alpha(1+\beta) + 2x\gamma(\alpha + \gamma - 1) + rx(\beta + \gamma)^2. \end{aligned}$$

and in this setting, G is negative if and only if the ratio is strictly larger than $1+r$. Now we have reformulated the game. The algorithm wants to have G equal to zero and can choose α and r for this purpose. In fact we can think of the algorithm wanting to maximize G , and choosing the smallest r such this goal is achieved for G being 0. In contrast the adversary wants to have G negative, and we can think of the adversary wanting to minimize G . For this purpose it can choose β, γ .

We analyze this game by fixing α and r and understanding the best response of the adversary. Since G is a quadratic expression in β and γ this best response can be either the minimizer of G , i.e. the value which sets the derivative of G to zero, or one of the boundaries of the domains of β and γ . Once we fixed the best response of the adversary we can analyze the best response of the algorithm to the choice of the adversary. This leads to a simple but tedious case analysis which

we present now. We start by breaking the analysis into two cases depending on how β compares to $2 + 1/p$, because this decides if the extreme point of G in β is a local minimum or maximum.

6.1. Case $x \geq 2 + 1/p$

For fixed α, r what would be the best response by the adversary? Note that G is convex both in β and γ . Hence one possibility for the adversary would be to choose the extreme points, provided that they satisfy the required bounds $0 \leq \beta^* \leq \alpha$ and $0 \leq \gamma^* \leq 1 - \alpha$. The extreme point for β is

$$\beta^* = \frac{\alpha - rx\gamma}{1 + rx}.$$

Choosing $\beta = \beta^*$ we observe that G remains convex in γ , since the second derivative of G in γ is

$$2x \left(2 + \frac{r}{1 + rx} \right).$$

Hence we consider the extreme point for γ which is

$$\gamma^* = \frac{(1 + rx)(1 - \alpha) - r\alpha}{2 + r + 2rx}.$$

Choosing $\gamma = \gamma^*$ the expression G writes as

$$G = pr + \frac{-2\alpha(r + 2) - (1 - \alpha)^2rx^2 + x(\alpha(\alpha + 2r - 2) + 1)}{2 + r + 2rx}.$$

We observe that G is concave in α hence one possibility for the algorithm is to choose the extreme point for α , which is

$$\alpha^* = \frac{x + rx^2 - rx - r - 2}{x + rx^2}.$$

Choosing $\alpha = \alpha^*$ the expression G writes as

$$\begin{aligned} G &= \frac{2 + r + (pr - 2)x + r(pr - 1)x^2}{x + rx^2} \\ &= \frac{2 - 2x + r(1 + px - x^2) + r^2px^2}{x + rx^2}. \end{aligned}$$

The algorithm chooses r such that G becomes zero. For this purpose we consider the roots of the numerator of G , which are

$$\begin{aligned} r_1^* &= \frac{x^2 - px - 1 - \sqrt{\Delta'}}{2px^2} \\ r_2^* &= \frac{x^2 - px - 1 + \sqrt{\Delta'}}{2px^2}, \end{aligned}$$

for

$$\Delta' = 8p(x - 1)x^2 + (1 + px - x^2)^2. \quad (9)$$

By case assumption $x \geq 2 + 1/p$, both roots of G are real. Since $G \geq 0$ means that the algorithm has ratio at most $1 + r$, and since the numerator of G is concave in r the algorithm has to choose the larger of both roots. Both roots are real, since $x \geq 1$ implies $\Delta' \geq 0$. We choose the root r_2^* which defines the ratio of the game, provided the following conditions are satisfied:

$$r \geq 0, 0 \leq \beta^* \leq \alpha^*, 0 \leq \gamma^* \leq 1 - \alpha^*.$$

Condition $r \geq 0$. The expression r_2^* has a single root in x namely at $x = 1$. Moreover its limit when x tends to infinity is $1/p$. Hence $r \geq 0$ holds by case assumption $x \geq 2 + 1/p > 1$.

Condition $\beta \geq 0$. With the choices for α, γ and r the value β^* writes as

$$\beta^* = \frac{-x^2 - px - 1 + \sqrt{\Delta'}}{2x^2} = r_2^*p - 1.$$

This means that the condition $\beta \geq 0$ translates into $r_2^* \geq 1/p$. We observe that $r_2^* = 1/p$ has a single solution in x , namely $x = 2 + 1/p$. Since at $x = 1$ the value of r_2^* is smaller than $1/p$, we know that for $x \geq 2 + 1/p$ we have $r_2^* \geq 1/p$ and hence $\beta \geq 0$.

Condition $\beta \leq \alpha$. We study the difference, which is

$$\alpha^* - \beta^* = \frac{1 + (-4 + p)x + 3x^2 - \sqrt{\Delta'}}{x^2 - 2x}.$$

There is no risk of dividing by 0 in the range $x \geq 2 + 1/p$. We observe that the numerator has two roots in x , namely $x = 0$ and $x = 1$. Moreover the numerator evaluates as $(2 + 4p)/p^2$ at $x = 2 + 1/p$. We conclude that $\beta^* \leq \alpha^*$ for the range $x \geq 2 + 1/p$.

Condition $0 \leq \gamma$. The expression γ writes as

$$\frac{-(x-1)^2 - px + \sqrt{\Delta'}}{x^3 - 2x^2},$$

again there is no risk of dividing by 0 in the range $x \geq 2 + 1/p$. The numerator evaluates to zero at $x = 0, x = 1$ and at $x = 1/(1+p)$. Since γ^* has the value $p^2/(1+3p+2p^2) > 0$ at $x = 2 + 1/p$ we obtain that $\gamma^* \geq 0$ in the range $x \geq 2 + 1/p$.

Condition $\gamma \leq 1 - \alpha$. We observe that the expression $1 - \alpha - \gamma$ simplifies as $1/x$.

This concludes the verification of the conditions on r, α, β, γ and shows that for $x \geq 2 + 1/p$ the competitive ratio is r_2^* .

6.2. Case $x < 2 + 1/p$

Intuitively, when x is small enough, then the delay caused by a job test will not be compensated by its benefit on the schedule. Indeed, we have observed earlier that the extreme point β^* is negative when $x < 2 + 1/p$. Hence for this range of x , the adversary chooses $\beta = 0$. This means that all tested jobs will be short and the algorithm has no incentive to test jobs, and will just execute them untested. We justify this intuition by a formal analysis. For the choice $\beta = 0$, the expression G reads as

$$G = pr - 2\alpha + \alpha^2 + x\gamma(r\gamma - 2(1 - \alpha - \gamma)),$$

which is convex in γ . Hence the adversary will choose the extreme point in γ which is

$$\gamma^* = \frac{1 - \alpha}{2 + r}.$$

Note that this value satisfies $0 \leq \gamma^* \leq 1 - \alpha$ as required. With this choice of γ the goal function becomes

$$pr(2+r) - x(1-\alpha)^2 - (2+r)(2-\alpha)\alpha$$

where we multiplied the goal function by $2+r$, preserving its sign and simplifying the expression. The second derivative in α is

$$4 + 2r - 2x,$$

which could be positive or negative.

We start with the non-positive case, meaning $2+r \leq x$. In this case the algorithm chooses the extreme value, which is $\alpha = 1$. In other words the algorithm tests all jobs, which all happen to be short and the competitive ratio becomes $(1+p)/p = 1 + 1/p$. But $2+r = 3 + 1/p$, which contradicts the case assumptions $2+r \leq x < 2 + 1/p$.

Finally we continue with the positive case. Here the algorithm chooses the lower bound for α , which is 0, translating the above mentioned intuition that the algorithm will execute all jobs untested.

For $\alpha = 0$, the expression G simplifies as $pr - x/(2+r)$ which is set to 0 by

$$r = \sqrt{(p+x)/p} - 1.$$

This concludes the case analysis, and therefore the proof of the theorem. \square

7. Adaptive model

In this section we analyze two-phase algorithms in the adaptive model, as described in the introduction. This means that the algorithm first tests some jobs, then executes untested the remaining jobs.

It is convenient to illustrate the interaction between such an algorithm and the adversary by a walk on a grid as follows. The vertices of the grid consist of all points with coordinates (c, d) such that $0 \leq c, d$ and $c + d \leq n$. Cells (c, d) with $c + d = n$ are called *final* cells, and non-final cells (c, d) are connected to the cells $(c+1, d)$ and $(c, d+1)$. These arcs form a directed acyclic graph with root $(0, 0)$, which is the upper left corner of the grid, see Figure 5.

The walk starts at the root, and follows only down or right steps to adjacent cells. If the algorithm decides to test a job, the adversary can choose to respond with a long job, resulting in a right step, or with a short job, resulting in a down step. In this sense the adversarial strategy played against a fictive algorithm testing all jobs, translates into a path P connecting the root to some final cell.

The interpretation of this walk is that if at some moment the current position is (c, d) , then it means that the algorithm tested $c + d$ jobs, among which d were long and c were short. These

two integers are not enough to fully describe the cost of the schedule produced by the algorithm. Therefore we introduce an additional value, called the *test cost*. Given a path P from the root to a cell (c, d) , the test cost associated to this cell is defined as the value $e = \sum_{(c', d')} (n - c')$, where the sum is taken over all grid cells (c', d') visited by the walk P between the root and the cell (c, d) but excluding the final cell (c, d) . This test cost e represents the increase in cost due to the different tests done so far, since each test delays by one all $n - c'$ subsequent job executions, i.e. its rank is $n - c'$. We associate to the cell (c, d) of the path P a so called *stop ratio*

$$R(c, d, e) = \max_{0 \leq b < n-c-d} \frac{\text{ALG}(b, c, d, e)}{\text{OPT}(b, d)},$$

for

$$\text{ALG}(b, c, d, e) = p \frac{n(n+1)}{2} + e + x \frac{(n-c)(n-c+1) - (n-c-b)(n-c-b+1) + d(d+1)}{2}$$

$$\text{OPT}(b, d) = p \frac{n(n+1)}{2} + x \frac{(b+d)(b+d+1)}{2}.$$

Note that the stop ratio $R(c, d, e)$ is associated to the cell (c, d) but depends also on the fixed path P , since it determines the value e .

This is the ratio reached by the game if the algorithm decides to switch to the execution phase after $c + d$ tests. Here b represents the number of long executed jobs (i.e. Ex) in the execution phase and ALG , OPT are the costs of the respective schedules, see Figure 4 for illustration.

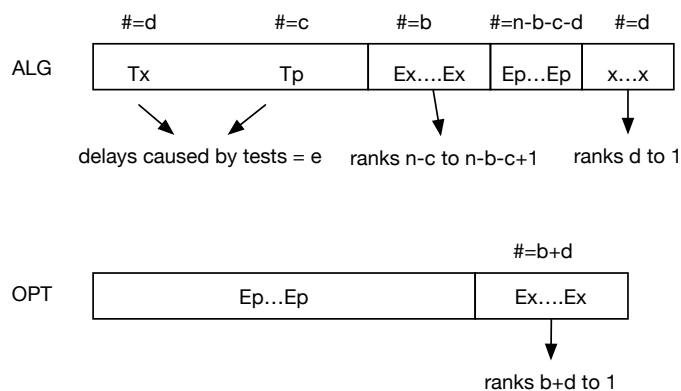


Figure 4: Illustration of the cost expression.

We explain now, how $R(c, d, e)$ can be computed efficiently.

Lemma 6. *The stop ratio $R(c, d, e)$ can be computed in constant time.*

Proof. The fraction $\text{ALG}(b, c, d, e)/\text{OPT}(b, d)$ is maximized over b , since the adversary gets to choose b and wants to maximize the ratio.

Both costs $ALG(b, c, d, e)$ and $OPT(b, d)$ are quadratic in b . Hence the ratio ALG/OPT as a function of b is continuous and derivable. We consider the function

$$g(b) = ALG(b, c, d, e)/OPT(b, d)$$

and try to identify a maximizer in $\{0, 1, \dots, n - c - d\}$. For this purpose we study the sign of $g(b) - g(b - 1)$. Since OPT is positive it is equivalent to analyze the sign of the function

$$\begin{aligned} h(b) &= \frac{2}{x}(ALG(b, c, d, e)OPT(b - 1, d) - ALG(b - 1, c, d, e)OPT(b, d)) \\ &= -2e(b + d) + pn(n + 1)(n - c - d - 2b + 1) \\ &\quad - x(b + d)(b(n - c + d + 1) - (1 + d)(n - c - d + 1)) \end{aligned}$$

where the purpose of the factor $2/x$ is to simplify the expression. Its second derivative in b is $-2x(n - c + d - 1) < 0$. Hence h is quadratic in b , and first increasing, then decreasing. This means that $g(b) - g(b - 1)$ is also first increasing in b , then decreasing. We distinguish different cases.

In case h is non-positive, we know that g is non-increasing, and hence maximized at $b = 0$. Otherwise h has two roots $b_1 < b_2$. At this point we consider $\lfloor b_2 \rfloor$. If it is negative or less than b_1 , then we know that h is non-positive at integral values $b \geq 0$, and again g is maximized at $b = 0$.

In the remaining case, $\lfloor b_2 \rfloor$ is non-negative and at least b_1 . Hence $\lfloor b_2 \rfloor$ is the last integer where h is non-negative and therefore is the integral maximizer of g . The roots are

$$b_{1,2} = \frac{(n - 2d^2 - c - d + 1)x - 2pn(n + 1) - 2e \pm \sqrt{\Delta}}{2x(n - c + d + 1)}$$

for

$$\begin{aligned} \Delta &= ((n - 2d^2 - c - d + 1)x - 2pn(n + 1) - 2e)^2 \\ &\quad + 4x(n - c + d + 1)((n - c - d + 1)(pn(n + 1) + xd(d + 1)) - 2de). \end{aligned}$$

We observe that

$$h(n - c - d) = -2e(n - c) - pn(n + 1)(n - c - d - 1) - x(n - c)(n - c + 1)(n - c - d - 1)$$

is negative, since $c + d < n$ is implied by Lemma 3. This implies $\lfloor b_2 \rfloor \leq n - c - d$, as required.

In summary, if $\Delta \leq 0$ or $\lfloor b_2 \rfloor < \max\{0, b_1\}$, then g is maximized at $b = 0$, otherwise it is maximized at $b = \lfloor b_2 \rfloor$. As a result the stop ratio $R(c, d, e)$ can be computed in constant time. \square

Finally we explain how to determine the optimal algorithmic strategy.

Theorem 3. *Assuming Conjecture 1, the optimal strategy for the algorithm in the adaptive setting can be computed in time $O(n^3)$.*

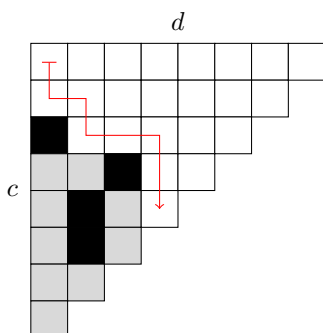


Figure 5: The grid as used in the procedure to compute the equilibrium schedule. Marked cells are black or gray. In red the boundary path.

Proof. We first show how to compute an optimal adversarial strategy, in form of a path P from the root to some final cell. Denote by R^* the minimal stop ratio along P . The algorithm will then walk along P , testing a new job at each step, until it reaches the cell with stop ratio R^* , at which point it switches to the execution phase. In case the adversary behaves differently than P , the procedure needs to restart and computes an optimal strategy for the adversary based on the current situation. For the ease of presentation we don't describe this generalized procedure.

Our algorithm to compute the optimal adversarial strategy works as follows. At any moment it stores a path P and its minimal stop ratio R^* and repeatedly tries to find a path with larger minimal stop ratio. In the formal description of Algorithm 3 we choose to represent P as a list of tuples (c, d, e) , where (c, d) is a grid cell and e the associated test cost.

Some cells of the grid are marked with the meaning that any path traversing any of the marked cells is guaranteed to have a minimal stop ratio at most R^* . The marked cells form a *combinatorial tableau* in the sense, that if cell (c, d) is marked, then so are all cells (c', d') with $c' \geq c$ and $d' \leq d$.

The procedure stops when the root cell $(0, 0)$ is marked and returns P as the optimal adversarial strategy.

For now assume that the root cell is not marked. Then there is a path P' which follows the boundary of the marked cells. This path has the property that for every cell (c, d) on that path, the associated test cost e is maximal among all paths that don't traverse a marked cell and therefore maximizes the stop ratio at (c, d) . Indeed, if one replaces two steps $(c', d') \rightarrow (c', d' + 1) \rightarrow (c' + 1, d' + 1)$ in a path by $(c', d') \rightarrow (c' + 1, d') \rightarrow (c' + 1, d' + 1)$ then the test cost of the cell $(c' + 1, d' + 1)$ is increased by 1 and so are the test costs of the subsequent cells on the path. The only path which does not allow this modification is the boundary path P' .

Two things can happen with P' . Either it contains a cell with stop ratio at most R^* . In what case this cell can be marked. We can safely mark as well all cells to its left and below, because any boundary path traversing these cells must traverse (c, d) as well, even if more cells get marked

later on.

Or the path has minimal stop ratio R' strictly larger than R^* . In what case P' is selected as the current best known adversarial strategy and R^* is increased to R' . As a result the cell on P' with stop ratio R^* will be marked in the next iteration as well. Hence a measure of progress of this procedure is the number of marked cells, which limits the number of iterations to $O(n^2)$, and probably much less in practice.

Inspecting the boundary path takes $O(n)$ time. The total time spend on marking cells is $O(n^2)$. Indeed when marking cells to the left and below a cell (c, d) on the boundary path, we know that cell $(c+1, d-1)$ is already marked, and it suffices to mark all cells $(c, d-1), (c, d-2), \dots$, until the boundary of the grid or a marked cell is reached and to mark all cells $(c+1, d), (c+2, d), \dots$, until the boundary of the grid or a marked cell is reached. As a result, the time spend on marking cells is constant for each cell. This means that the overall time complexity is $O(n^3)$. \square

8. Experiments

We conducted experiments to verify Conjecture 1, for up to 10 jobs, 128 uniformly spread values² $p \in (0, 100]$, as well as 128 uniformly spread values $x \in (0, 10]$, and for both the adaptive and non-adaptive model. By analyzing the game tree of each instance, no counterexample to Conjecture 1 was found. The programs used for the experiments are provided in the companion webpage <https://www.lip6.fr/Christoph.Durr/Scheduling-w-Oracle/>. Note that these experiments explore the game tree which has 2^{2n} leafs, so $n = 10$ is roughly the limit up to which we can test the conjecture. The algorithm from the previous section is based on the two-phase assumption and has theoretical complexity $O(n^3)$, which in practice is $O(n^2)$ because it seems to terminate after a constant number of iterations. Hence an implementation of this algorithm can clearly handle roughly 10 000 jobs.

Just for curiosity we plot the algorithmic strategies in terms of number of tested jobs, see Figures 6 and 7. It is interesting to observe that the strategies are quite different in the adaptive and in the non-adaptive model. We also plot the competitive ratio, and not surprisingly observe that it is worse when x is large and p small, while tending to 1 for larger values of p . Another measure that is interesting to extract from these experiments, is the gain of adaptivity. It is defined as the competitive ratio in the non-adaptive model compared to its counterpart in the adaptive model, and has been studied theoretically in contexts of query algorithms [11, 24]. We can observe that the problem has a small gain of adaptivity, in the order of 2%.

²We used a power of two in the hope of reducing rounding errors. In addition we conducted experiments using the Python module *fractions* which completely avoids these errors.

Algorithm 3 Algorithm computing the optimal strategy in the adaptive model

```
procedure BOUNDARYPATH( $R^*$ , grid)
  ( $c, d, e$ ) = (0, 0, 0)
  Initially  $P$  contains only (0, 0, 0)
  while  $c + d < n$  do
    if ( $c + 1, d$ ) is not marked then ▷ privilege down steps
       $c = c + 1$ 
    else
       $d = d + 1$ 
    end if
    if  $R(c, d) \leq R^*$  then ▷ boundary path does not improve  $R^*$ 
      mark all cells ( $c', d'$ ) with  $c' \geq c, d' \leq d$ 
      return None
    end if
    append ( $c, d, e$ ) to  $P$ 
     $e = e + c$ 
  end while
  return  $P$ 
end procedure

procedure OPTIMALADVERSARY
  grid =  $n \times n$  grid without marked cells
   $R^* = 1$ 
  while cell (0,0) is not marked do
     $P = \text{BoundaryPath}(R^*, \text{grid})$ 
    if  $P$  is not None then ▷  $P$  improves  $R^*$ 
       $R^* = \text{minimum stop ratio along } P$ 
      best = ( $P, R^*$ ) ▷ best path-ratio pair
    end if
  end while
  return best
end procedure

procedure OPTIMALALGORITHM
  ( $P, R^*$ ) = OptimalAdversary()
  ( $c, d, e$ ) = (0, 0, 0)
  while  $R(c, d, e) > R^*$  do
    Test next job, execute immediately if it is short
    ( $c, d, e$ ) is next cell on  $P$  ▷ assumes adversary follows optimal strategy  $P$ .
  end while ▷ stop at minimum stop ratio
  Execute untested all remaining jobs, followed by the  $d$  postponed tested long jobs
end procedure
```

We know that assuming Conjecture 1, in the non-adaptive model, all equilibrium schedules are of the form $(Tx)^*(Tp)^*(Ex)^*(Ep)^*$. We observe that in the adaptive model for most instances the equilibrium schedule is also of this form. Figure 6 shows instances where this is not the case. For larger value of n , we observed quite a variety of these *noticeable schedules*, as we call them. They do not seem to obey a particular structure, which could be exploited in a more time efficient algorithm for the adaptive model.

9. Conclusion

This paper studies a single machine scheduling problem with the objective of minimizing the total completion times. The novelty is that every job can either be short or long, but this information is initially hidden to the algorithm. The scheduler can execute a processing oracle for a given job in order to learn this information, but jobs can also be scheduled untested. The performance of an algorithm is compared with the objective value of the optimal schedule, which benefits on full knowledge of the processing times. It is conjectured that an optimal scheduler follows a two-phase strategy. In the first phase some jobs are tested, and scheduled right away if they turn out to be short. In the second phase all untested jobs are executed. First we show how to produce the optimal non-adaptive strategy in time $O(n^2)$, where n is the number of given jobs. Then we compute its asymptotic competitive ratio. For the adaptive model, we describe how to compute the optimal strategy in time $O(n^3)$, assuming the above mentioned conjecture. Finally we conduct various experiments on this problem, measuring the competitive ratio, the number of tests and the gain of adaptivity as a function of the job length parameters p and x .

We leave Conjecture 1 as an open problem. Future research directions include the improvement of the running times, for example by a subtle use of binary search. In addition the next step could be the study of a more general problem, where every job j is known to have a processing time in a given interval $[\underline{p}_j, \bar{p}_j]$ and has a testing time q_j . And finally, randomization clearly helps against the oblivious adversary, already by initially shuffling the given jobs. Hence it would be interesting to analyze the randomized competitive ratio.

Acknowledgment

We would like to thank Thomas Erlebach for helpful discussions as well as several anonymous referees for helpful comments on a previous version of this manuscript. This work was partially supported by the French research agency (ANR-18-CE25-0008 and ANR-19-CE48-0016), by the EPSRC grant EP/S033483/1 as well as by University of Santiago of Chile (USACH), Beca Olga Ulianova, ANID, Beca de Postdoctorado en el Extranjero, Folio: 74200020 and ANID, Proyecto FONDECYT, Folio: 1211640.


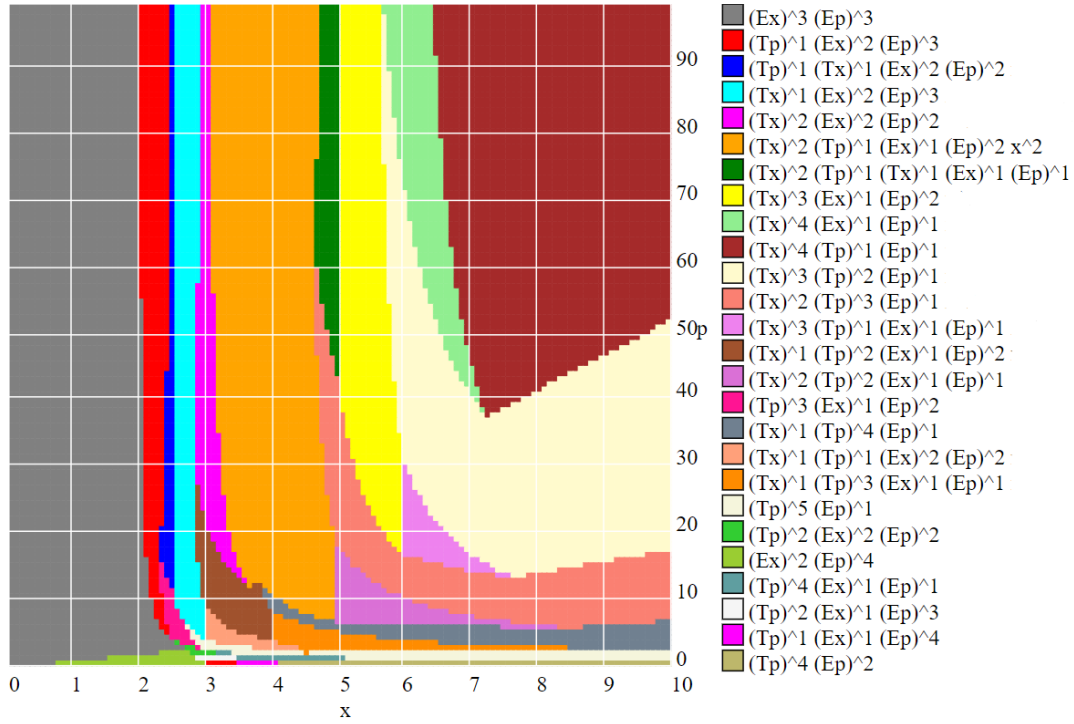


Figure 6: Experiments made for $n = 6$ jobs. Every (x, p) point corresponds to the game with specific values for x and p . For each point the equilibrium schedule is computed, and depicted in the first plot. In this example two equilibrium schedules differ from the pattern $(Tx)^*(Tp)^*(Ex)^*(Ep)^*$, namely $TpTxExExEpEp$, and $TxTxTpTxExEp$. Second plot shows the competitive ratio with the color scheme documented to its right.




Figure 7: More experiments made for $n = 6$ jobs. The first plot is a simplification of the first plot of Figure 6, showing only the number of tests done by the algorithm in the equilibrium schedule. The second plot compares the situation with the non-adaptive setting. It is interesting to observe that the region of points (x, p) for which the equilibrium schedule has the same fixed number of tests seems to be connected in the non-adaptive case, while it is not in the adaptivity case. The third plot shows the gain of adaptivity, with a different scale for p for improved readability. It is interesting to see that it is not monotone in p nor in x .

References

- [1] José García and Alvaro Peña. Robust optimization: concepts and applications. *Nature-inspired Methods for Stochastic, Robust and Dynamic Optimization*, page 7, 2018. Publisher: BoD–Books on Demand.
- [2] Bram L Gorissen, Ihsan Yanıkođlu, and Dick den Hertog. A practical guide to robust optimization. *Omega*, 53:124–137, 2015. Publisher: Elsevier.
- [3] Kurt Marti. *Stochastic optimization methods*. Springer, 2nd edition, 2015.
- [4] Simon Kahan. A model for data in motion. In *Proceedings of the twenty-third annual ACM symposium on Theory of computing - STOC '91*, pages 265–277, New Orleans, Louisiana, United States, 1991. ACM Press.
- [5] Retsev Levi. Scheduling (Talk at Dagstuhl Seminar 16081). In *Dagstuhl Reports*, volume 6. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016. Issue: 2.
- [6] Retsev Levi, Thomas Magnanti, and Yaron Shaposhnik. Scheduling with Testing. *Management Science*, 65(2), 2018.
- [7] Christoph Dürr, Thomas Erlebach, Nicole Megow, and Julie Meißner. Scheduling with Explorable Uncertainty. In *The 9th Innovations in Theoretical Computer Science Conference (ITCS)*, 2018.
- [8] Christoph Dürr, Thomas Erlebach, Nicole Megow, and Julie Meißner. An Adversarial Model for Scheduling with Testing. *Algorithmica*, pages 3630–3675, 2020.
- [9] Susanne Albers and Alexander Eckl. Explorable Uncertainty in Scheduling with Non-Uniform Testing Times. In *Proc. of the 18th Workshop on Approximation and Online Algorithms*, September 2020.
- [10] Tomas Feder, Rajeev Motwani, Rina Panigrahy, Chris Olston, and Jennifer Widom. Computing the median with uncertainty. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 602–607. ACM, 2000.
- [11] Anupam Gupta, Viswanath Nagarajan, and Sahil Singla. Adaptivity gaps for stochastic probing: Submodular and XOS functions. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1688–1702. SIAM, 2017.
- [12] Steven Chaplick, Magnús M. Halldórsson, Murilo S. de Lima, and Tigran Tonoyan. Query Minimization under Stochastic Uncertainty. In Y. Kohayakawa and F. K. Miyazawa, editors, *LATIN 2020: 14th Latin American Theoretical Informatics Symposium*, volume 12118 of *Lecture Notes in Computer Science*, pages 181–193. Springer Berlin Heidelberg, 2020.

- [13] Tomás Feder, Rajeev Motwani, Liadan O’Callaghan, Chris Olston, and Rina Panigrahy. Computing shortest paths with uncertainty. *Journal of Algorithms*, 62(1):1–18, 2007. Publisher: Elsevier.
- [14] Richard Bruce, Michael Hoffmann, Danny Krizanc, and Rajeev Raman. Efficient update strategies for geometric computing with uncertainty. *Theory of Computing Systems*, 38(4):411–423, 2005. Publisher: Springer.
- [15] Thomas Erlebach, Michael Hoffmann, Danny Krizanc, Matúš Mihalák, and Rajeev Raman. Computing Minimum Spanning Trees with Uncertainty. In *STACS 2008*, pages 277–288. IBFI Schloss Dagstuhl, 2008.
- [16] Jacob Focke, Nicole Megow, and Julie Meißner. Minimum Spanning Tree under Explorable Uncertainty in Theory and Experiments. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 75. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- [17] Nicole Megow, Julie Meißner, and Martin Skutella. Randomization helps computing a minimum spanning tree under uncertainty. *SIAM Journal on Computing*, 46(4):1217–1240, 2017. Publisher: SIAM.
- [18] Thomas Erlebach, Michael Hoffmann, and Frank Kammer. Query-competitive algorithms for cheapest set problems under uncertainty. *Theoretical Computer Science*, 613:51–64, 2016. Publisher: Elsevier.
- [19] Thomas Erlebach. Computing and Scheduling with Explorable Uncertainty. In *Conference on Computability in Europe*, pages 156–160. Springer, 2018.
- [20] Yaron Shaposhnik. *Exploration vs. exploitation: reducing uncertainty in operational problems*. PhD Thesis, Massachusetts Institute of Technology, 2016.
- [21] Daniel D. Sleator and Robert E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, February 1985.
- [22] Susanne Albers. Online algorithms: a survey. *Mathematical Programming*, 97(1):3–26, July 2003.
- [23] Dennis Komm. *An Introduction to Online Computation: Determinism, Randomization, Advice*. Texts in Theoretical Computer Science. An EATCS Series. Springer International Publishing, 2016.
- [24] Brian C Dean, Michel X Goemans, and Jan Vondrák. Adaptivity and approximation for stochastic packing problems. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 395–404. Society for Industrial and Applied Mathematics, 2005.