



HAL
open science

Interruptible algorithms for multiproblem solving

Spyros Angelopoulos, Alejandro Lopez-Ortiz

► **To cite this version:**

Spyros Angelopoulos, Alejandro Lopez-Ortiz. Interruptible algorithms for multiproblem solving. *Journal of Scheduling*, 2020, 23 (4), pp.451-464. 10.1007/s10951-020-00644-9 . hal-02986029

HAL Id: hal-02986029

<https://hal.science/hal-02986029>

Submitted on 2 Nov 2020

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.

Interruptible Algorithms for Multiproblem Solving

Spyros Angelopoulos*

Alejandro López-Ortiz

October 29, 2018

Abstract

In this paper we address the problem of designing an interruptible system in a setting in which n problem instances, all equally important, must be solved concurrently. The system involves scheduling executions of contract algorithms (which offer a trade-off between allowable computation time and quality of the solution) in m identical parallel processors. When an interruption occurs, the system must report a solution to each of the n problem instances. The quality of this output is then compared to the best-possible algorithm that has foreknowledge of the interruption time and must, likewise, produce solutions to all n problem instances. This extends the well-studied setting in which only one problem instance is queried at interruption time.

In this work we first introduce new measures for evaluating the performance of interruptible systems in this setting. In particular, we propose the *deficiency* of a schedule as a performance measure that meets the requirements of the problem at hand. We then present a schedule whose performance we prove that is within a small factor from optimal in the general, multiprocessor setting. We also show several lower bounds on the deficiency of schedules on a single processor. More precisely, we prove a general lower bound of $(n + 1)/n$, an improved lower bound for the two-problem setting ($n = 2$), and a tight lower bound for the class of round-robin schedules. Our techniques can also yield a simpler, alternative proof of the main result of Bernstein *et al.* [4] concerning the performance of cyclic schedules in multiprocessor environments.

Keywords: Anytime computation; contract algorithms; interruptible algorithms; acceleration ratio; scheduling problems in Artificial Intelligence; performance measures in scheduling.

1 Introduction

A designer of real-time systems should anticipate the situation in which there are limitations on the available execution time. Applications such as medical diagnosis systems, automated trading systems and game-playing programs require that the system may be queried at any time during its execution, at which point a solution must be reported. *Anytime algorithms* occur precisely in such settings, namely in situations in which a computationally difficult problem is addressed under uncertain running time availability. Such algorithms will produce an output whose quality improves as a function of the available computation time. Anytime algorithms were introduced by Horvitz [11], [12] and Dean and Boddy [7] and arise in central AI problems such as heuristic search, and planning under uncertainty [18].

*Sorbonne Université, CNRS, Laboratoire d'informatique de Paris 6, LIP6, F-75252 Paris, France.
spyros.angelopoulos@lip6.fr

Russel and Zilberstein [16] distinguish between two main classes of anytime algorithms. On the one hand, the class of *interruptible algorithms* consists of algorithms that can be interrupted at any point during their execution, and must always report their current (albeit not necessarily optimal) solution. On the other hand, the class of *contract algorithms* consists of algorithms which specify the exact amount of allowable computation time as part of their input. Such algorithms must terminate their execution before a solution can be produced, otherwise the output may be meaningless.

Interruptible algorithms offer, by their definition, more flexibility; in contrast, contract algorithms are typically much easier to design, implement and maintain [5]. Thus, it is desirable, in general, to be able to transform a contract algorithm to its interruptible variant. This can be addressed in an algorithm-specific manner, but we are interested, instead, in “black-box” techniques that are algorithm-independent. Towards this end, a well-studied technique is by scheduling executions of contract algorithms, in a machine that consists either of a single or even multiple processors.

More precisely, in the most general setting, we are presented with a set P of n problem instances which we want to solve, and for each problem we are given a contract algorithm for the said problem. In addition, we are given a system of m identical processors on which we can schedule this sequence of contract algorithms. The goal is to devise an efficient schedule, that is a strategy that assigns interleaved executions of all contract algorithms on the processors. In the standard setting, upon an interruption, a query for a problem in P is issued. The algorithm must then report the solution of the (completed) contract algorithm with the longest execution time for the queried problem, since the latter provides the best completed solution by the time of interruption. Thus, we would like these “lengths” or durations of completed contracts to be as large as possible, since the longer the execution time, the better the quality of the solution returned by the contract algorithm (and thus by the interruptible system as well).

The standard performance measure for a schedule of contract algorithms is the *acceleration ratio* [16]. Informally, the measure describes the multiplicative increase in processor speed required for the schedule to compensate for the lack of knowledge of the interruption time. More formally, let $l_{p,t}$ denote the length of the largest contract for problem p completed by time t in a schedule X . The acceleration ratio $\alpha(X)$ of schedule X is then defined as

$$\alpha(X) = \sup_t \max_{p \in P} \frac{t}{l_{p,t}}. \tag{1}$$

Thus, the acceleration ratio is a worst-case measure that compares the quality of the solution returned by the schedule (that is, the quantity $l_{p,t}$) to an ideal, optimal algorithm that knows the interruption t in advance and dedicates a single processor in order to run a contract of length t for problem p . It is worth noting that for $n = m = 1$, i.e., for the setting of a single problem and a single processor, the problem of devising a schedule of minimum acceleration ratio is known in the context of online computation as the *online bidding* problem [6] (though, to our knowledge, previous work has not identified this equivalence).

Related work Simulating interruptible algorithms by means of schedules of contract algorithms has been a topic of extensive study. Russell and Zilberstein [16] were the first to present such an explicit simulation. For the case of one problem instance and a single processor, they provided a schedule based on iterative doubling of contract lengths for which the corresponding interruptible

algorithm has acceleration ratio at most four. Zilberstein *et al.* [19] showed that in this case the schedule is optimal, in the sense that no other schedule of better acceleration ratio exists.

Zilberstein *et al.* [19] addressed the generalization of multiple problem instances (assuming a single available processor), and Bernstein *et al.* [5] studied the generalization in which contracts for a single problem instance must be scheduled on a set of multiple processors. For both cases, optimal schedules are derived. Bernstein *et al.* [4] addressed the problem in its full generality, namely the setting in which n problem instances are given and the schedule is implemented on m processors. In particular, they showed an upper bound of $\frac{n}{m} \left(\frac{m+n}{n}\right)^{\frac{m+n}{m}}$ on the acceleration ratio; in addition, they showed that the schedule is optimal for the class of *cyclic* schedules. The latter is a somewhat restricted, but still very rich and intuitive class of schedules with round-robin characteristics. This restriction was removed by López-Ortiz *et al.* [14], who showed that this acceleration ratio is indeed optimal among all possible schedules. Angelopoulos *et al.* [3] studied the setting in which the interruptions are not absolute deadlines, but instead an additional window of computational time may be provided.

The problem of devising schedules of optimal acceleration ratio has interesting parallels with another well studied problem in both AI and Operations Research, namely the problem of searching for a hidden target in an environment that consists of a number of unbounded, concurrent lines: this is known as the *star search* or *ray search* problem. Bernstein *et al.* [4] were the first to establish connections between the two problems; more recently, [2] explored further connections between these classes of problems in probabilistic, fault-tolerant and randomized settings. In both works, the acceleration ratio of the scheduling strategy is compared to the *competitive ratio*, which is the standard performance measure of search strategies.

Contribution The central observation that motivates our work is that the acceleration ratio becomes problematic, as a performance measure, when at interruption time the algorithm is required to return a solution to *all n problems in P* instead of only to a specific queried problem. This arises, for instance, in systems which involve parallel executions of different heuristics (and at interruption time, the best heuristic is chosen). Another example is a medical diagnostic system which must perform concurrent evaluations for a number of medical issues. Here, the decision of the expert may very well have to take into account all such evaluations.

Suppose that we are given a set of n problems (indexed $0, \dots, n-1$) and a set M of m identical processors (indexed $0, \dots, m-1$). Suppose, in particular, that $n > m$, and for the purposes of illustrating our argument, say that $n \gg m$. Note that it is not feasible for any ideal, optimal algorithm (that is, an algorithm with advance knowledge of the interruption time t) to schedule n contracts of length t to m processors, since $n > m$. In a sense, if we applied the acceleration ratio to this domain, we would compare the performance of an interruptible algorithm which is expected to make progress on all n problems to an algorithm which only makes optimal progress on at most m : such a comparison is rather not fair. This shortcoming was noticed by Zilberstein *et al.* [19], who define a scaled-down variant of the acceleration ratio for the case of n problems and one processor as the quantity $\sup_t \max_{p \in P} \frac{t/n}{l_{p,t}}$. This measure describes, informally, an even distribution of the processor time among the n problem instances for the optimal (offline) schedule.

A different way of arguing about the above shortcoming is to consider the scenario in which at interruption time t a single problem $p \in P$ is queried. Naturally, the interruptible algorithm does not know neither t or p in advance, but then the optimal offline algorithm should be oblivious of p as well, otherwise it would make progress only on p while ignoring all other problems in P .

This suggests that the offline optimal algorithm implicit in the definition of the acceleration ratio is overly powerful, and better measures are needed for the setting that we study.

Motivated by the above observations, in this paper we address the problem of designing interruptible algorithms using schedules of executions of contract algorithms, assuming that m identical processors are available, and n problem instances, all equally significant, must be solved. We begin by considering measures alternative to the acceleration ratio (Section 2), and we propose the *deficiency* of a schedule as our measure of choice. In Section 3 we present a schedule whose deficiency is very small and rapidly decreasing in the ratio n/m (in contrast, the acceleration ratio of every schedule is known to approach infinity when $n/m \rightarrow \infty$). More precisely, we show analytically that its deficiency is at most 3.74 if $n \leq m$, and at most 4, if $n > m$. A numerical evaluation provides even smaller values.

Even though the deficiency of the proposed schedule is small, we provide further theoretical justification for our choice of schedule. Namely, in Section 4 we present several lower bounds on the deficiency of schedules in the single-processor setting ($m = 1$). More precisely, we prove a general lower bound of $(n + 1)/n$, an improved bound for the two-problem setting ($n = 2$), and a tight lower bound for round-robin schedules. We also remark that the schedule is optimal for the setting $n = m = 1$. The proofs are based on techniques originally developed in the context of search theory [1] that have also been applied to scheduling problems [14, 3], and which allow us to relate the performance of schedules with arbitrary contract lengths to that of schedules with exponentially increasing lengths. As a further illustration of the applicability of these techniques, we give a much simpler, alternative proof of the main result in [4], namely we identify the optimal acceleration ratio that can be achieved by cyclic schedules, in the multi-processor setting.

As a last remark for this section, it is worth noting that revisiting the definition of an ideal, optimal algorithm (and introducing new measures that capture this weakening) is an often encountered concept. As an illustrative example, a multitude of measures alternative to the competitive ratio have been introduced in the context of the analysis of online algorithms (see e.g., the survey [8]). These measures were motivated by the observation that the concept of an offline optimal algorithm is often overly powerful, with the undesirable side-effect that many online algorithms are often deemed theoretically optimal, though very inefficient in practice. As a second example, we note that for the problem of multi-target searching in a star, [15, 13] consider a weakening of the optimal algorithm which knows the distance of the targets from the original position of the searcher, but not the precise ray on which each target is located. In both examples, the introduction of new measures gives rise to new algorithmic ideas and new techniques for analysis.

2 Problem formulation and comparison of measures

Consider an (infinite) schedule X of executions of contract algorithms (also called *contracts*). For an interruption time t we denote by $l(X, p, t)$ the length of the longest contract for problem $p \in P$ which is finished by time t in X (or simply $l_{p,t}$ when X is implied from context). We make the canonical assumption that at interruption time at least one contract per problem has already completed its execution.

We need to formalize the question: what is the best way to exploit the available resources (i.e., processors), in order to solve the set of problems P ? Towards this end, consider a schedule Y , which, in contrast to X , is finite: more precisely, Y schedules n distinct contracts, one for each problem in P (if Y schedules more than one contract per problem, then we can transform Y to a

schedule Y' which is at least as good as Y by keeping only the largest contract per problem that appears in Y). Each contract in Y is scheduled in one of the m processors in M . We require that Y is *feasible with respect to t* , in the sense that the *makespan* of Y , namely the total sum of contract lengths on the most loaded processor used by Y does not exceed t . Let \mathcal{Y}_t denote the class of all schedules Y with the above properties. We will be calling Y an *offline solution* since it relies on advance knowledge of t .

Having defined \mathcal{Y}_t , we need a measure of how a schedule $Y \in \mathcal{Y}_t$ compares to X , which will also dictate which is the *best* schedule in \mathcal{Y}_t compared to X . First observe that the acceleration ratio is not an appropriate measure for our setting, since under it the optimal offline schedule Y for interruption t dedicates all its resources to the contract that is worked on the least by X while failing to produce an answer for all other problems. This results in a large acceleration ratio which however does not truly reflect the quality of X (effectively, the optimal solution “cheats” by ignoring all but one problem instances, which is not acceptable in our setting).

Another alternative would be to compare the smallest contract completed by X to the smallest contract completed by Y , by time t . We will need some preliminary definitions first. Let S_X^t denote the set $\{l(X, j, t) | j \in [0, n - 1]\}$ (that is, the set of the largest contracts per problem completed by time t) and let $S_X^t(i)$ denote the i -th smallest element of S_X^t . Similarly, for $Y \in \mathcal{Y}_t$ let S_Y^t denote the set of n contracts in Y , and $S_Y^t(i)$ be the i -th smallest contract in Y , respectively (ties are resolved arbitrarily).

Formally, we define the *performance ratio of X with respect to Y at time t* as:

$$\text{perf}(X, Y, t) = \frac{\min_i S_Y^t(i)}{\min_i S_X^t(i)} = \frac{S_Y^t(1)}{S_X^t(1)}. \quad (2)$$

The performance ratio of X at time t is then defined as

$$\text{perf}(X, t) = \sup_Y \text{perf}(X, Y, t),$$

where Y is a feasible schedule wrt t . Last, the performance ratio of X is defined as $\text{perf}(X) = \sup_t \text{perf}(X, t)$.

The first observation is that under this measure there exists an “optimal” offline schedule in which all contracts have the same length: here, by “optimal” we mean a schedule $Y \in \mathcal{Y}_t$ against which the performance ratio of X is maximized. Indeed, given any offline schedule Y , consider any schedule Y' such that the length of all its contracts is $\min_i S_Y^t(i)$. Note that such a feasible Y' exists, since all contracts in Y are at least that long, and Y itself is feasible. Then it follows from the definition that $\text{perf}(X, Y, t) = \text{perf}(X, Y', t)$. We can show a close relationship of the performance ratio to the acceleration ratio for the standard setting (namely, when we seek the solution only to the queried problem).

Lemma 1. *There is a schedule X such that for any arbitrary interruption time t*

$$\text{perf}(X, t) = \begin{cases} \frac{n}{m} \left(\frac{m+n}{n}\right)^{\frac{m+n}{m}} & \text{for } m \geq n \\ \frac{n}{m} \frac{1}{\lceil n/m \rceil} \left(\frac{m+n}{n}\right)^{\frac{m+n}{m}} & \text{for } m < n. \end{cases}$$

Furthermore, X is optimal with respect to this measure.

Proof. Consider first the case $m \geq n$, then an optimal offline schedule consists of executing one contract of length t per problem, each on its own processor and hence $\text{perf}(X, t) = \frac{t}{\min_i S_X(i)}$. Note then that $\text{perf}(X) = \sup_t \text{perf}(X, t)$ is precisely the definition of the acceleration ratio $\alpha(X)$, for which the results of [4] and [14] show the optimal value of $\frac{n}{m} \left(\frac{m+n}{n}\right)^{\frac{m+n}{n}}$.

Next, suppose that $m < n$, then every offline schedule Y is such that there exists at least one processor in which at least $\lceil n/m \rceil$ contracts are scheduled. As argued earlier, there exists an optimal offline schedule in which all contracts have the same length. It follows then that there is an optimal offline schedule with contract lengths equal to $t/\lceil n/m \rceil$. Therefore, $\text{perf}(X, t) = \frac{t/\lceil n/m \rceil}{\min_i S_X(i)}$ and $\text{perf}(X) = \sup_t \text{perf}(X, t)$ can then be described as $\frac{1}{\lceil n/m \rceil} \cdot \alpha(X)$, and the lemma follows. \square

We note that for the case of one processor ($m = 1$) and n problems, Lemma 1 shows that the performance ratio of the optimal schedule matches the measure proposed by [19], as mentioned in Section 1.

It is also not difficult to show that, unlike the optimal acceleration ratio (which is bounded by a function linear in $\frac{n}{m}$), the optimal performance ratio is bounded by a small constant; in other words, the schedule in the proof of Lemma 1 is very efficient under this refined measure.

Proposition 1. *The performance ratio of the schedule in the proof of Lemma 1 is bounded by $2e$, if $n > m$, and by 4 , if $n \leq m$.*

Proof. In the case $m \geq n$, we have

$$\text{perf}(X) = \left(1 + \frac{n}{m}\right) \left(1 + \frac{m}{n}\right)^{\frac{n}{m}} \leq 2 \cdot 2 = 4.$$

In the case $m < n$ we have that

$$\text{perf}(X) \leq \left(1 + \frac{m}{n}\right) \left(1 + \frac{n}{m}\right)^{\frac{n}{m}} \leq 2 \cdot e = 2e.$$

In either case, $\text{perf}(X) \leq 2e$. \square

Figure 1 shows the plot of $\text{perf}(X)$ as a function of the ratio $\frac{n}{m}$ for $n > m$, and assuming, for simplicity, that m divides n . Note how the performance of the schedule is between 4 and e , and decreases rapidly as n/m increases. When $m \geq n$, $\text{perf}(X)$ decreases in a similar manner, as m/n increases, and takes values between 4 and 1.

While the performance ratio seems a better candidate for a measure in the context of our problem than the acceleration ratio, it is far from being the best possible choice. Note that according to this measure, every contract in the optimal offline solution may have fixed length, namely $t/\lceil n/m \rceil$. While it is guaranteed that the smallest contract in S_X^t indeed does not exceed $t/\lceil n/m \rceil$, the solution produced by X may be such that there exist several contracts in S_X^t which exceed this length. More formally, there may exist j such that $S_X^t(j) > S_Y^t(j)$ for the optimal offline solution Y . This is clearly undesirable, since it becomes difficult to argue that the optimal offline solution Y at time t is indeed better than the solution produced by the interruptible algorithm at time t , even though, supposedly, Y is optimal.

The above motivates the need for defining a further measure, one which takes into account the intuitive expectation that the optimal offline solution should perform better than the interruptible algorithm *on each problem*. To accomplish this, we allow the offline solution to observe the behavior




Figure 1: Plot of function $\text{perf}(X)$, assuming m divides n .

of X at each point in time t , and then produce an appropriate “optimal” offline solution, tailored to the specifications of our problem. In a sense we allow the offline solution sufficient power in choosing its schedule, while at the same time we require that it produces solutions to all problem instances. This yields a measure which we call *deficiency*, and which is: i) consistent with the requirements of the problem; and ii) powerful enough, in the sense that if an algorithm performs well with respect to the new measure, then there are very strict guarantees about its performance.

Formally, we say that Y is *at least as good as* X , denoted by $Y \geq X$, if and only if $S_Y^t(i) \geq S_X^t(i)$, for all $i \in [1, n]$. Then for a schedule $Y \in \mathcal{Y}_t$ with $Y \geq X$, we say that the *deficiency of X wrt Y* for interruption t is defined as

$$\text{def}(X, Y, t) = \min_i \frac{S_Y^t(i)}{S_X^t(i)} \quad (3)$$

The *deficiency of X given t* is then defined as

$$\text{def}(X, t) = \sup_{Y \in \mathcal{Y}_t, Y \geq X} \text{def}(X, Y, t) \quad (4)$$

We define the deficiency of X simply as

$$\text{def}(X) = \sup_t \text{def}(X, t) \quad (5)$$

While (3) gives a formal definition of the deficiency, we will obtain an alternative definition that provides us with more flexibility in terms of the analysis of actual schedules. To this end, we will relate this measure to the *makespan* of a schedule.

Definition 1. *Suppose that we are given m identical processors and a set S of n jobs, where each job has a certain size (or length). For a given schedule of S (i.e., for an assignment of S to the processors), the makespan of the schedule is the load of the least loaded processor in this schedule, where the load of a processor is defined as the sum of the sizes of the jobs assigned to the said processor. We denote by $OPT(S)$ the optimal makespan of S , among all possible schedules.*

We will first show how to select a canonical representative from the set of schedules of optimal deficiency.

Lemma 2. *Given schedule X and interruption t , there exists an optimal schedule Y with the following properties:*

1. Y is such that $S_Y^t(i) = d \cdot S_X^t(i)$, for all $1 \leq i \leq n$ and $d = \text{def}(X, t) \geq 1$.
2. d is the largest possible value such that S_Y^t can be feasibly scheduled in M . i.e., the makespan of the corresponding schedule is exactly t .

Proof. Let Y' be an optimal schedule given X and t . It is not hard to transform Y' to another optimal schedule Y which satisfies property (1) as follows. First, observe that from the definition of $\text{def}(X, t)$ it follows that $S_{Y'}^t(i) \geq d \cdot S_X^t(i)$. Now for all i such that $S_{Y'}^t(i) > d \cdot S_X^t(i)$, we can decrease the length of each contract $S_{Y'}^t(i)$ obtaining a new schedule Y for which $S_Y^t(i) = d \cdot S_X^t(i)$. Clearly, Y is a feasible optimal schedule as well, for the given X . For property (2) assume otherwise, i.e. that the makespan of Y is not t . Since Y is feasible with respect to t , by definition we have that the makespan of Y must equal $t - \epsilon$ for some $\epsilon > 0$. Now consider a new schedule Y'' identical to Y except that every contract length is multiplied by $t/(t - \epsilon)$. The makespan of Y'' is t , and thus the deficiency of X is at least

$$\text{def}(X, Y'', t) = \min_i \{S_{Y''}(i)/S_X(i)\} = dt/(t - \epsilon),$$

which contradicts the fact that $\text{def}(X, t) = d$. □

Lemma 2 implies the following corollary which establishes the relation between the deficiency and the makespan.

Corollary 1. $\text{def}(X, t) = \frac{t}{OPT(S_X^t)}$, where $OPT(S_X^t)$ is the minimum makespan for scheduling n jobs in m processors, with job i having size (length) equal to $S_X^t(i)$.

Similarly to the acceleration ratio, one can think of the deficiency of a schedule as a resource-augmentation measure. Specifically, it guarantees that if the contracts of X are scheduled with a processor speedup of $\text{def}(X)$, then for any interruption t , no schedule that knows t can obtain a better feasible solution relatively to the one achieved by X .

3 A near-optimal schedule for general m

In this section we propose a schedule for which we will show that the deficiency is bounded by a small constant, and is, thus, very close to optimal. More precisely, we will identify an efficient schedule that belongs in the class of *exponential* schedules. Informally, these are schedules in which problems are assigned to processors in round-robin order, and in which the lengths of the assigned contracts increases geometrically.

Definition 2 ([4]). A schedule X with contract lengths¹ x_0, x_1, \dots is called round-robin if it satisfies the following properties:

1. Problem round-robin: For every i , the contract x_i is assigned to problem $i \bmod n$.
2. Processor round-robin: For every i , the contract x_i is assigned to processor $i \bmod m$.

¹With a slight abuse of notation, given a sequence of contract lengths x_0, x_1, \dots , we will often refer to the contract of length x_i as the contract x_i . This is only done for simplicity, and we do not require that contracts have pairwise different lengths.

Definition 3. A round-robin schedule X with contract lengths x_0, x_1, \dots is called exponential if the contract length x_i is equal to b^i , for some $b > 1$.

Since an exponential schedule is fully described by its base, the remainder of this section is devoted to finding a value b that yields a schedule of small deficiency. Following [4], we will denote by G_k the finish time of the k -th contract of the schedule, in the round-robin order, whereas $L_k = b^k$ denotes the length of the k -th contract in this order.

The following lemma shows that in order to evaluate the deficiency of any schedule (not necessarily exponential) it suffices to consider only interruption times right before a contract terminates. Let G_c^- denote a time infinitesimally smaller than the finish time G_c of contract $c \in X$. Recall also that according to Definition 1, $OPT(S)$ denotes the optimal makespan for a schedule of a set S of jobs in m processors.

Lemma 3. $\text{def}(X) = \sup_{c \in X} \frac{G_c^-}{OPT(S_X^{G_c^-})}$.

Proof. Let t_1, t_2, \dots be the sequence of finish times of contracts in X , in increasing order. For any t such that $t_i < t < t_{i+1} = G_c$ (for some contract c) we have that $S_X^t = S_X^{G_c^-}$ (since no contract finishes in (t_i, t_{i+1})). From Corollary 1 we have

$$\text{def}(X, t) = \frac{t}{OPT(S_X^t)} = \frac{t}{OPT(S_X^{G_c^-})} \leq \frac{G_c^-}{OPT(S_X^{G_c^-})}.$$

The lemma follows from $\text{def}(X) = \max_t \text{def}(X, t)$. □

Since X is a round-robin, exponential schedule, Lemma 3 implies that

$$\text{def}(X) = \sup_{k \geq 0} \frac{G_{n+k}^-}{OPT(\{L_k, \dots, L_{n+k-1}\})} \quad (6)$$

where $L_i = b^i$ is the length of the i -th contract in the round-robin order. Eq. (6) suggests that in order to bound the deficiency of X , one needs to obtain a lower bound on the makespan of n jobs, whose lengths are equal to L_k, \dots, L_{n+k-1} , respectively. This is accomplished in the next lemma.

Lemma 4. For $L_i = b^i$ it holds that

$$OPT(L_k, \dots, L_{n+k-1}) \geq \kappa \cdot b^k \frac{b^{n+m-1} - b^{(n-1) \bmod m}}{b^m - 1},$$

where κ is defined as $\kappa = \max \left\{ \frac{1}{2^{-1/m}}, \frac{b^m - 1}{b^m} \right\}$.

Proof. We will evaluate the makespan of the well-known greedy algorithm which schedules jobs of sizes L_k, \dots, L_{n+k-1} , in m identical processors numbered $0, \dots, m-1$, assuming the jobs are considered in this particular order (i.e., in increasing order of size). More precisely, the algorithm will assign each job to the machine of least current load. It is easy to see that the decisions of the greedy algorithm are such that the job of size L_{k+i} is scheduled on processor $i \bmod m$. Moreover, the incurred makespan is determined by the total load of jobs scheduled on the same processor as

job L_{n+k-1} , namely processor $(n-1) \bmod m$. Denote by $Gr(L_k, \dots, L_{n+k-1})$ the makespan of the greedy algorithm. We obtain

$$\begin{aligned}
Gr(L_k, \dots, L_{n+k-1}) &= \sum_{i=0}^{\lfloor (n-1)/m \rfloor} b^{k+mi+(n-1) \bmod m} \\
&= b^k b^{(n-1) \bmod m} \sum_{i=0}^{\lfloor (n-1)/m \rfloor} b^{mi} \\
&= b^k \frac{b^{n+m-1} - b^{(n-1) \bmod m}}{b^m - 1}.
\end{aligned} \tag{7}$$

To complete the proof, we need to argue that the greedy scheduling policy does not achieve a makespan worse than $(1/\kappa)$ times the optimal makespan. Graham's fundamental theorem on the performance of the greedy scheduling policy [10] states that the greedy algorithm has an approximation ratio of $2 - 1/m$. Moreover, we know that the optimal makespan is at least b^{n+k-1} , which in conjunction with (7) yields that the greedy algorithm is also a $b^m/(b^m - 1)$ approximation (for our specific instance). The lemma follows by combining the above two approximation guarantees. \square

We now proceed to bound the deficiency of the exponential schedule X . It is easy to show that for exponential schedules, $G_{n+k} = \frac{b^{k+n+m} - b^{(k+n) \bmod m}}{b^m - 1}$ [4]. Let $\lambda = 1/\kappa = \min \left\{ 2 - \frac{1}{m}, \frac{b^m}{b^m - 1} \right\}$. Combining with (6) and Lemma 4 we obtain

$$\begin{aligned}
\text{def}(X) &\leq \lambda \cdot \sup_{k \geq 0} \frac{b^{k+n+m} - b^{(k+n) \bmod m}}{b^k (b^{n+m-1} - b^{(n-1) \bmod m})} \\
&= \lambda \cdot \sup_{k \geq 0} \frac{b^{n+m} - (b^{(k+n) \bmod m})/b^k}{b^{n+m-1} - b^{(n-1) \bmod m}} \\
&\leq \lambda \cdot \frac{b^{n+m}}{b^{n+m-1} - b^\gamma},
\end{aligned} \tag{8}$$

where γ is defined as $(n-1) \bmod m$.

We thus seek the value of b that minimizes the RHS of (8). It is not clear that this can be done analytically, since the factor λ depends on b , and the derivative of the RHS does not have roots that can be identified analytically. Instead, let $f(b)$ denote the function $b^{n+m}/(b^{n+m-1} - b^\gamma)$. Then $f(b)$ is minimized for a value of b equal to $\beta = (n+m-\gamma)^{\frac{1}{n+m+\gamma-1}}$. Let ρ be such that $n-1 = \rho m + \gamma$, then $\beta = (m(\rho+1)+1)^{\frac{1}{m(\rho+1)}}$. Observe also that $f(b) = 1/(b^{-1} - b^{\gamma-n-m}) = 1/(b^{-1} - b^{-m(\rho+1)-1})$. Summarizing, for $b = \beta$ we obtain a schedule of deficiency

$$\text{def}(X) \leq \min \left\{ 2 - \frac{1}{m}, \frac{\beta^m}{\beta^m - 1} \right\} \cdot \frac{1}{\beta^{-1} - \beta^{-m(\rho+1)-1}}. \tag{9}$$

Figure 2 shows a plot of the deficiency of the schedule (or, more accurately, the RHS of (9) for the interesting case where $n > m$) as function of m and ρ . Note that for fixed m , the deficiency increases as a function of n , until a point at which it becomes relatively stable with n (this can be explained by the factor λ that is the minimum of two functions, one of which does not depend on n). For large m , and even larger n , the deficiency is close to 2. From the plot, the maximum value of deficiency is $3/8 \cdot 5^{5/4} \approx 2.803$. Notwithstanding the numerical results, we can bound analytically the deficiency of the schedule X , as shown in the following proposition.




Figure 2: Plot of $\text{def}(X)$ (as bounded by (9)) as a function of m and ρ .

Proposition 2. *For the exponential schedule X in which $b = \beta$, we have that $\text{def}(X) \leq 3.74$, if $n > m$, and $\text{def}(X) \leq 4$, if $n \leq m$.*

Proof. Substituting $m(\rho + 1)$ with y in (9) we obtain

$$\text{def}(X) \leq \left(2 - \frac{1}{m}\right) \frac{1}{(y+1)^{-\frac{1}{y}} - (y+1)^{-\frac{y+1}{y}}}. \quad (10)$$

Using standard calculus it is straightforward to show that the functions $(y+1)^{-1/y}$ and $(y+1)^{-\frac{y+1}{y}}$ are increasing and decreasing functions of y , respectively. Therefore, the denominator of the RHS of (10) is an increasing function of $y = m(\rho + 1)$. It turns out that for $n > m$ this upper bound on $\text{def}(X)$ is maximized when $m = 2$ and $\rho = 1$, hence $\text{def}(X) \leq 3.74$. When $n \leq m$, the $(2 - 1/m)$ factor vanishes, since the greedy schedule achieves optimal makespan (in other words, $\lambda = 1$). In this case $\text{def}(X) \leq 4$. \square

We conclude this section by observing that the exponential schedule X we obtained outperforms the strategy of optimal acceleration ratio, in all ranges of n and m . More precisely, the schedule of [4] and [14] has unbounded deficiency for $m \gg n$. For $n > m$, this schedule has constant deficiency, but larger than the deficiency of X (namely, a deficiency equal to 4.24 in the worst case).

4 Lower bounds on schedule deficiency for a single processor

In this section we show several lower bounds on the deficiency of schedules, assuming $m = 1$. In Section 4.1 we show a general lower bound on the deficiency of a schedule for n problems. In Section 4.2 we give a matching lower bound for the class of round-robin schedules. The result shows that the exponential schedule of Section 3 is optimal for this class of schedules. Last, in Section 4.3 we give an improved lower bound on the deficiency of a schedule for two problems on a single processor. The proof demonstrates, in addition, some of the difficulties that one has to bypass in order to improve the lower bound on arbitrary schedules, and will hopefully provide some intuition about how to handle the general case.

We first note that, for $m = 1$, we obtain from (8) that an exponential schedule X with base b has deficiency

$$\text{def}(X) \leq \frac{b^{n+1}}{b^n - 1}, \quad (11)$$

which is minimized for a value of b equal to $(n + 1)^{\frac{1}{n}}$.

Corollary 2. *The deficiency of the best exponential schedule for n problems is at most $\frac{(n+1)^{\frac{n+1}{n}}}{n}$.*

In what follows, we will denote by EXP the best exponential schedule whose deficiency is described by Corollary 2.

Recall that we denote by $l_{p,t}$ the length of the longest contract for problem p that has completed by time t . We observe that for all schedules X on a single processor,

$$\text{def}(X) = \sup_t \frac{t}{\sum_{i=0}^{n-1} l_{i,t}} = \sup_{t \in T_X} \frac{t}{\sum_{i=0}^{n-1} l_{i,t}}, \quad (12)$$

where T_X is defined as the set of all times right before the completion of a contract in X .

4.1 A general lower bound

Theorem 1. *For any schedule X that involves n problems and a single processor, we have that*

$$\text{def}(X) \geq \frac{n+1}{n}.$$

Proof. Suppose, by way of contradiction, that $\text{def}(X) < \frac{n+1}{n}$. From (12), for any given time t we have

$$\frac{n+1}{n} > \text{def}(X) \geq \frac{t}{\sum_{i=0}^{n-1} l_{i,t}},$$

therefore $\sum_{i=0}^{n-1} l_{i,t} > t \frac{n}{n+1}$. This implies there exists a problem for which its largest contract completed by time t has length at least $\frac{t}{n+1}$. Let $q \leq t$ denote the completion time of this contract. From the definition of the deficiency, and by the observation that $\sum_{i=0}^{n-1} l_{i,q^-} \leq q - \frac{t}{n+1}$, we obtain that

$$\text{def}(X) \geq \frac{q}{q - \frac{t}{n+1}} \geq \frac{t}{t - \frac{t}{n+1}},$$

where the last inequality follows from the monotonicity of the function $f(x) = x/(x - a)$, with $a > 0$. Therefore, $\text{def}(X) \geq \frac{n+1}{n}$, a contradiction. \square

Figure 3 illustrates the performance of the best exponential schedule vs the lower-bound value of Theorem 1, for $n = 1 \dots 20$.




Figure 3: Plot of the deficiency of the best exponential schedule (in red) vs the the lower-bound value of Theorem 1 (in blue).

4.2 Round-robin schedules on a single processor

We now consider the class of all round-robin schedules on a single processor. Namely, a schedule X in this class schedules contracts of lengths x_0, x_1, \dots , in this order, such that the contract x_i is assigned to problem $i \bmod n$. We can assume, without loss of generality, than for all $i \geq n$, $x_i < x_{i+n}$, otherwise the contract x_{i+n} can be omitted from the schedule without increasing the deficiency of the schedule. We will show that the exponential schedule EXP is optimal for this class of schedules. To this end, we will apply a theorem that allows us to relate the performance of schedules that have certain structure (such as round-robin schedules) to the performance of exponential schedules.

More precisely, we will make use of the results by Gal [9] and Schuierer [17] which we state here in a simplified form. Given an infinite sequence $X = (x_0, x_1, \dots)$, define $X^{+i} = (x_i, x_{i+1}, \dots)$ as the suffix of the sequence X starting at x_i . Define also $G_a = (1, a, a^2, \dots)$ to be the geometric sequence in a .

Theorem 2 ([9, 17]). *Let $X = (x_0, x_1, \dots)$ be a sequence of positive numbers, r an integer, and $a = \limsup_{n \rightarrow \infty} (x_n)^{1/n}$, for $a \in \mathbb{R} \cup \{+\infty\}$. Let F_k , $k \geq 0$ be a sequence of functionals which satisfy the following properties:*

1. $F_k(X)$ only depends on x_0, x_1, \dots, x_{k+r} ,
2. $F_k(X)$ is continuous for all $x_i > 0$, with $0 \leq i \leq k+r$,
3. $F_k(\alpha X) = F_k(X)$, for all $\alpha > 0$,
4. $F_k(X + Y) \leq \max(F_k(X), F_k(Y))$, and
5. $F_{k+1}(X) \geq F_k(X^{i+1})$, for all $i \geq 1$,

then

$$\sup_{0 \leq k < \infty} F_k(X) \geq \sup_{0 \leq k < \infty} F_k(G_a).$$

At an intuitive level, Theorem 2 allows us to lower-bound the supremum of an infinite set of functionals by the supremum of simple functionals of a geometric sequence; the latter is typically

much easier to compute than the former. In our context, the objective is then to express the deficiency as the supremum of a sequence of functionals. This technique has been applied in search games [1], but also in previous work on scheduling of contract algorithms [14, 3].

Theorem 3. *For every round robin schedule X on a single processor, we have that $\text{def}(X) \geq \frac{(n+1)^{\frac{n+1}{n}}}{n}$.*

Proof. Consider a time t right before the completion of contract x_{k+n} of the round-robin schedule, with $k \geq 0$. The set S_X^t of the n largest contracts per problem is then the set $\{x_k, x_{k+1}, \dots, x_{k+n-1}\}$. Thus,

$$\text{def}(X) \geq \frac{t}{\sum_{j=0}^{n-1} x_{k+j}} = \frac{\sum_{j=0}^{k+n} x_j}{\sum_{j=k}^{k+n-1} x_j}.$$

Define now the functional $F_k(X) = \frac{\sum_{j=0}^{k+n} x_j}{\sum_{j=k}^{k+n-1} x_j}$. It is easy to verify that $F_k(X)$ satisfies the conditions of Theorem 2; one can also appeal to Example 7.3 in [1]. Therefore, we conclude that there is $a > 0$ such that

$$\text{def}(X) = \sup_{0 \leq k < \infty} F_k(X) \geq \sup_{0 \leq k < \infty} \frac{\sum_{j=0}^{k+n} a^j}{\sum_{j=k}^{k+n-1} a^j}. \quad (13)$$

Note that if $a = 1$, we have that $\frac{\sum_{j=0}^{k+n} a^j}{\sum_{j=k}^{k+n-1} a^j} = \frac{k+n+1}{n}$, which tends to infinity as $k \rightarrow \infty$, and thus $\text{def}(X) = \infty$ in this case. Thus we can assume that $a \neq 1$, thus (13) implies

$$\text{def}(X) \geq \sup_{0 \leq k < \infty} \frac{a^{k+n+1} - 1}{a^{k+n} - a^k} = \sup_{0 \leq k < \infty} \frac{a^{n+1} - \frac{1}{a^k}}{a^n - 1}. \quad (14)$$

Note that if $a < 1$, then the RfHS of (14) is ∞ . Thus, we can assume that $a > 1$. In this case,

$$\sup_{0 \leq k < \infty} \frac{a^{n+1} - \frac{1}{a^k}}{a^n - 1} = \frac{a^{n+1}}{a^n - 1}. \quad (15)$$

Last, the expression $\frac{a^{n+1}}{a^n - 1}$ attains its minimum at $a = (n+1)^{\frac{1}{n}}$. For this value, and combining (14) and (15) we obtain that

$$\text{def}(X) \geq \frac{(n+1)^{\frac{n+1}{n}}}{n},$$

which concludes the proof. \square

From Corollary 2, it follows that EXP is optimal for the class of round-robin schedules.

In order to illustrate the further applicability of Theorem 2, we will show how it can lead to a simple alternative proof of the main result in [4]. Namely, we will show that there is an exponential schedule for scheduling contracts for n problems in m parallel processors that achieves optimal *acceleration ratio* for the class of *cyclic* schedules. This class of schedules was introduced in [4] and consists of round-robin schedules (i.e., schedules that satisfy the properties of Definition 2), with the additional *length-increasing* property: namely, suppose that i, j are such that the contracts x_i, x_j are assigned to the same problem p . Then, if $i < j$, we require that $x_i < x_j$.

Benstein *et al.* [4] gave an exponential cyclic schedule with acceleration ratio $\frac{n}{m} \left(\frac{n+m}{n} \right)^{\frac{n+m}{m}}$. Their main result showed that this schedule is optimal for the class of cyclic schedules:

Theorem 4 ([4]). *For scheduling contracts for n problems in m processors, every cyclic schedule X has acceleration ratio $\alpha(X)$ such that*

$$\alpha(X) \geq \frac{n}{m} \left(\frac{n+m}{n} \right)^{\frac{n+m}{m}}.$$

Alternative proof of Theorem 4. Let $X = (x_0, x_1, \dots)$ denote a given cyclic schedule. For a given $k \geq 0$, let p_k denote the processor to which contract x_{k+n+m} is assigned, and let the set I_k denote the indices of contracts scheduled in p_k and completed up to, and including the completion time of x_{k+n+m} . We have

$$\alpha(X) \geq \max_{l \in [k, k+m-1]} \frac{\sum_{i \in I_l} x_i}{x_{l+m}}. \quad (16)$$

This is because the contract x_{l+n+m} is completed at time $\sum_{i \in I_l} x_i$, and because right before this time, the longest contract completed for problem l has length x_{l+m} . The latter follows from the length-increasing and problem-round-robin properties of the schedule.

Using the property $\max(a/c, b/d) \geq (a+b)/(c+d)$, for $a, b, c, d > 0$, (16) gives

$$\alpha(X) \geq \frac{\sum_{l=k}^{k+m-1} \sum_{i \in I_l} x_i}{\sum_{l=k}^{k+m-1} x_{l+m}} \quad (17)$$

Due to the processor-round-robin property of X , we have that

$$\sum_{l=k}^{k+m-1} \sum_{i \in I_l} x_i = \sum_{l=0}^{k+n+2m-1} x_l,$$

and since (17) holds for all $k \geq 0$, we obtain that

$$\alpha(X) \geq \sup_{0 \leq k < \infty} \frac{\sum_{l=0}^{k+n+2m-1} x_l}{\sum_{l=k+m}^{k+2m-1} x_l}.$$

Define now the functional $F_k(X) = \frac{\sum_{l=0}^{k+n+2m-1} x_l}{\sum_{l=k+m}^{k+2m-1} x_l}$. Once again, it is easy to verify that $F_k(X)$ satisfies the conditions of Theorem 2. Therefore, we conclude that there is $a > 0$ such that

$$\alpha(X) \geq \sup_{0 \leq k < \infty} \frac{\sum_{l=0}^{k+n+2m-1} a^l}{\sum_{l=k+m}^{k+2m-1} a^l}, \text{ for some } a \geq 0. \quad (18)$$

Similar to the proof of Theorem 3, it is easy to see that if $a \leq 1$, then the RHS of (18) is ∞ . We can thus assume that $a > 1$, in which case we obtain that

$$\begin{aligned} \alpha(X) &\geq \sup_{0 \leq k < \infty} \frac{a^{k+n+2m} - 1}{a^{k+m}(a^m - 1)} \\ &= \sup_{0 \leq k < \infty} \frac{a^{n+m} - \frac{1}{a^{k+m}}}{a^m - 1} = \frac{a^{n+m}}{a^m - 1}. \end{aligned} \quad (19)$$

Last, the expression $\frac{a^{n+m}}{a^m - 1}$ attains its minimum at $a = ((m+n)/n)^{1/m}$. For this value of a , (19) gives

$$\alpha(X) \geq \frac{n}{m} \left(\frac{n+m}{n} \right)^{\frac{n+m}{m}},$$

which concludes the proof. \square

4.3 Schedule deficiency for $n \in \{1, 2\}$, $m = 1$

In this section we will show improved lower bounds on the deficiency of schedules that involve two problems and a single processor. First, we argue that for $n = 1$, the exponential schedule EXP is optimal. Second, we show how to obtain a lower bound on the deficiency of a schedule for $n = 2$ that improves upon the lower bound of Theorem 1. Namely, we obtain a lower bound of 2.115 whereas Theorem 1 gives a lower bound of only 1.5. Note that the corresponding upper bound due to EXP is 2.598.

We can assume, without loss of generality, that for any given problem p , if an optimal schedule begins the execution of two contracts for problem p at times t_1, t_2 , with $t_1 < t_2$, then the length of the contract starting at time t_1 is strictly smaller than the length of the contract starting at time t_2 (otherwise, the contract starting at time t_2 may be omitted altogether from the schedule, without increasing its deficiency).

Proposition 3. *For $n = 1$ and $m = 1$, EXP has optimal deficiency.*

Proof. Let $X = (x_0, x_1, \dots)$ denote a given schedule. From (12), we have

$$\text{def}(X) = \sup_t \max_p \frac{t}{l_{p,t}}.$$

Thus, the deficiency of X is identical to its acceleration ratio $\alpha(X)$. From [16], we know that the optimal acceleration ratio equals 4, and is achieved by an exponential schedule which is precisely X . \square

Next, we consider the setting $n = 2$, $m = 1$. We will first prove the existence of optimal schedules with a useful property. Informally, the property states that there exists an optimal schedule such that any time a new contract is about to be scheduled, the problem that has been worked the least will be chosen. We will call such optimal schedules *normalized*. The next lemma proves this property for general n , assuming a single processor.

Lemma 5. *For scheduling contracts for n problems in a single processor, there exists a schedule of optimal deficiency which satisfies the following property: If at time T a new contract is about to be started, then it is assigned to a problem i which minimizes $l_{i,T}$.*

Proof. Let X denote an optimal schedule, and suppose, by way of contradiction, that T is the first time in which a contract is about to begin its execution in X that does not satisfy the property. More precisely, suppose that at time T , X executes a contract for a problem $j \neq i$ such that $l_{j,T} > l_{i,T}$. We will show that there exists a schedule X' such that X' is also an optimal schedule, and it satisfies the property for all $t \leq T$. By successively applying a series of such transformations, we obtain an optimal schedule that satisfies the property.

We define a schedule X' which is identical to X , with the exception that all contracts in X which are assigned to problems i and j (defined as above) and which start their execution at time $t > T$ will switch problem assignment. In other words, a contract that was assigned to problem i will now be assigned to problem j , and vice versa. Let T' be the first time such that $T' > T$ and a contract for problem i is executed in X . Moreover, let L and L' denote the lengths of the two contracts for problems j and i that begin their executions at times T and T' , respectively, in X . From the definition of X and (12), we observe that for all t such that $t < T + L$ or $t \geq T' + L'$, we have that $\text{def}(X', t) = \text{def}(X, t)$.

It thus remains to show that $\text{def}(X', t) \leq \text{def}(X, t)$ for all $t \in [T + L, T' + L')$. For any such t , we have that

$$\text{def}(X, t) = \frac{t}{\sum_{k \notin \{i, j\}} l_{k,t} + l_{j,t} + l_{i,T}},$$

whereas

$$\text{def}(X', t) = \frac{t}{\sum_{k \notin \{i, j\}} l_{k,t} + l_{j,t} + l_{j,T}}.$$

Since $l_{i,T} < l_{j,T}$, it follows that $\text{def}(X', t) < \text{def}(X, t)$, for all $t \in [T + L, T' + L')$, which concludes the proof. \square

Lemma 6. *For $n = 2$ and $m = 1$ there is an optimal normalized schedule in which at most two contracts per problem are scheduled consecutively.*

Proof. Let X denote a normalized schedule. Let t denote the earliest time in X such that two consecutive contracts, say for problem $p = 0$, are about to be scheduled, and let x_i, x_{i+1} denote their lengths respectively ($x_{i+1} > x_i$). Let also l_0, l_1 , denote the largest contracts for problems 0, 1, respectively, that are completed in X by time t , and λ_1 the largest contract for problem 1 completed by time t^- , i.e., infinitesimally smaller than t . Since X is normalized, we have that $\lambda_1 \leq l_0 \leq l_1$, and $x_i \leq l_1$.

Consider the schedule X' that is obtained by X by omitting the contract x_{i+1} . Namely, all other contracts in X' are scheduled in the same order, with no idle time (see Figure 4 for an illustration). Note that X' remains normalized. We will show that at least one of the following properties hold: i) $\text{def}(X') \leq \text{def}(X)$; or ii) $x_{i+1} \geq l_1$. If the first property holds, this implies that we can effectively remove the first of the successively scheduled contracts without affecting the optimality of the schedule. If the second property holds, since X is normalized, this means that the contract to be scheduled right after x_{i+1} in X will be a contract for problem 1. The combination of these observations implies the lemma.




Figure 4: An illustration of schedules X (top) and X' (bottom). X' is derived by effectively removing x_i from X .

We thus must compare the deficiency X to that of X' . To this end, note that for X one needs only to consider interruptions that occur at t^- , as well as right before contracts x_i, x_{i+1} are completed, whereas for X' , one needs to consider interruptions at time t^- and right before contract

x_{i+1} is completed (in X'). All other interruptions have the same, or better contribution to the deficiency in X' than in X . Define Q_X and $Q_{X'}$ as

$$\begin{aligned} Q_X &= \max\left(\frac{t}{l_0 + \lambda_1}, \frac{t + x_i}{l_0 + l_1}, \frac{t + x_i + x_{i+1}}{x_i + l_1}\right) \text{ and} \\ Q_{X'} &= \max\left(\frac{t}{l_0 + \lambda_1}, \frac{t + x_{i+1}}{l_0 + l_1}\right), \end{aligned}$$

then it follows that if $Q_{X'} \leq Q_X$ then $\text{def}(X') \leq \text{def}(X)$.

First, we observe that if $\frac{t}{l_0 + \lambda_1} \geq \frac{t + x_{i+1}}{l_0 + l_1}$, then $Q_{X'} \leq Q_X$ and thus $\text{def}(X') \leq \text{def}(X)$. Thus we may assume that $\frac{t}{l_0 + \lambda_1} < \frac{t + x_{i+1}}{l_0 + l_1}$. Combined with $\lambda_1 \leq l_0$, this implies that $t < \frac{2l_0}{l_1 - l_0} x_{i+1}$, therefore we obtain

$$t + x_{i+1} < \frac{l_0 + l_1}{l_1 - l_0} x_{i+1}. \quad (20)$$

We also observe that if $\frac{t + x_i + x_{i+1}}{x_i + l_1} \geq \frac{t + x_{i+1}}{l_0 + l_1}$, then again $Q_{X'} \leq Q_X$ and thus $\text{def}(X') \leq \text{def}(X)$. Thus we may also assume that $\frac{t + x_i + x_{i+1}}{x_i + l_1} < \frac{t + x_{i+1}}{l_0 + l_1}$, from which we obtain that

$$t + x_{i+1} > \frac{x_i(l_0 + l_1)}{x_i - l_0} > \frac{l_1(l_0 + l_1)}{l_1 - l_0}. \quad (21)$$

From (20) and (21) it follows that

$$\frac{l_1(l_0 + l_1)}{l_1 - l_0} < t + x_{i+1} < \frac{l_0 + l_1}{l_1 - l_0} x_{i+1},$$

which in turn implies that $x_{i+1} > l_1$. This concludes the proof. \square

We will now use the property of the optimal schedules shown in Lemma 6 in combination with Theorem 2 so as to obtain an improved lower bound on the deficiency of any schedule.

Theorem 5. *For $n = 2$ and $m = 1$, any schedule has deficiency at least 2.115.*

Proof. From Lemma 6, there exists a schedule X of optimal deficiency such that at most two consecutive contracts are executed for each problem. Let $X = (x_0, x_1, \dots)$ denote such a schedule. Consider an interruption right before contract x_{k+1} is completed, and suppose, without loss of generality, that the said contract is for problem 0. Then, from Lemma 6, it follows that either x_k and x_{k-1} are the largest contracts for problems 0 and 1, respectively, that are completed by the interruption, or x_k and x_{k-2} are the largest such contracts for problems 1 and 0, respectively. Therefore, the deficiency of X is at least

$$\text{def}(X) \geq \sup_{k \geq 0} \frac{\sum_{j=0}^{k+1} x_j}{x_k + x_{k-1} + x_{k-2}}.$$

Define now the functional $F_k(X) = \frac{\sum_{j=0}^{k+1} x_j}{x_k + x_{k-1} + x_{k-2}}$. Once again, it is easy to verify that $F_k(X)$ satisfies the conditions of Theorem 2. Therefore, we conclude that there is $a > 0$ such that

$$\text{def}(X) \geq \sup_{k \geq 0} \frac{\sum_{j=0}^{k+1} a^j}{a^k + a^{k-1} + a^{k-2}}, \text{ for some } a \geq 0. \quad (22)$$

Similar to the proof of Theorem 3, it is easy to see that if $a < 1$, then the LHS of (22) is ∞ . We can thus assume that $a > 1$, in which case we obtain that

$$\begin{aligned} \text{def}(X) &\geq \sup_{k \geq 0} \frac{\sum_{j=0}^{k+1} a^j}{a^k + a^{k-1} + a^{k-2}} \\ &= \sup_{k \geq 0} \frac{a^4 - 1/a^{k-2}}{(a-1)(1+a+a^2)} = \frac{a^4}{a^3-1}. \end{aligned}$$

Last, the expression $\frac{a^4}{a^3-1}$ attains its minimum at $a = 2^{2/3}$. For this value of a , we have that $\text{def}(X) \geq 2.115$. \square

5 Conclusion

In this paper we introduced and studied the problem of devising interruptible algorithms by means of schedules of contract algorithms, in a setting in which solutions to all problem instances are required and all problems are equally important. This generalizes the case where only one problem is queried at interruption time, and for which the standard performance measure is the acceleration ratio. We introduced the deficiency of a schedule as a measure that reflects performance in a more faithful manner than either the acceleration ratio or the performance ratio. We next presented a schedule whose deficiency we showed is bounded by a small constant (at most 3.74 if $n \leq m$ and at most 4, if $n > m$; numerical evaluation provides even smaller values). Furthermore, for the case of one processor, we presented several lower bounds on the deficiency of a schedule in a variety of settings, assuming a single processor.

Even though our proposed solutions are efficient, it would be nevertheless interesting to know whether our schedule is theoretically optimal for any number of processors. This appears to be a challenging task. Even for the case of a single processor ($m = 1$) and n problems, the deficiency depends on the sum of n contract lengths, and this quantity appears in the denominator of the fractions that describe the deficiency (12). This creates several technical complications that make the application of tools such as Theorem 2 quite difficult. As shown in Lemma 6, it is possible that successive contracts for the same problem may be executed, which implies that Theorem 2 cannot yield an optimal lower bound, unless one derives further structural properties of the optimal schedule.

The situation becomes even more complicated for general m , since minimizing makespan is NP-hard. More importantly, one needs an upper bound on the makespan of a schedule (i.e., the output of a makespan-minimization algorithm) in a closed-form expression as a function of the sizes of jobs. Graham's greedy algorithm yields such a closed-form expression, but it is no better than a 2-approximation. Thus, unlike the case $m = 1$, it is not clear how to express the exact deficiency even of simple schedules such as the exponential schedule.

Another direction for further research is to consider the generalization in which at interruption time a subset of $i \leq n$ problems are queried for their solution. We are also interested in a related setting in the context of robot searching in multiple concurrent rays. [4] and [2] showed interesting connections between ray searching for a single target under the competitive ratio and contract scheduling under the acceleration ratio. Suppose, however, that instead of a single target, i targets must be located, at unknown positions from the common origin. Suppose we have m robots available to explore the rays. What is the best strategy for the robots so as to locate the targets? There are

interesting parallels between the two problems, e.g., the optimal solution in the multi-target variant can be formulated as a makespan scheduling problem, with job sizes a function of the distances of the targets from the origin. We believe the approach in this paper could be useful in addressing this problem.

References

- [1] S. Alpern and S. Gal. *The Theory of Search Games and Rendezvous*. Kluwer Academic Publishers, 2003.
- [2] S. Angelopoulos. Further connections between contract-scheduling and ray-searching problems. In *Proceedings of the 24th International Joint Conference in Artificial Intelligence (IJCAI)*, pages 1516–1522, 2015.
- [3] S. Angelopoulos, A. López-Ortiz, and A. Hamel. Optimal scheduling of contract algorithms with soft deadlines. *Journal of Scheduling*, 20(3):267–277, 2017.
- [4] Daniel S. Bernstein, Lev Finkelstein, and Shlomo Zilberstein. Contract algorithms and robots on rays: Unifying two scheduling problems. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1211–1217, 2003.
- [5] D.S. Bernstein, T. J. Perkins, S. Zilberstein, and L. Finkelstein. Scheduling contract algorithms on multiple processors. In *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI)*, pages 702–706, 2002.
- [6] M. Chrobak and C. Kenyon-Mathieu. Competitiveness via doubling. pages 115–126, 2006.
- [7] T. Dean and M. S. Boddy. An analysis of time-dependent planning. In *Proceedings of the 15th National Conference on Artificial Intelligence (AAAI)*, pages 49–54, 1998.
- [8] R. Dorrigiv and A. López-Ortiz. A survey of performance measures for on-line algorithms. *SIGACT News (ACM Special Interest Group on Automata and Computability Theory)*, 36(3):67–81, 2005.
- [9] S. Gal. *Search Games*. Academic Press, 1980.
- [10] R. Graham. Bounds for certain microprocessing anomalies. *Bell system Tech. J.*, 45:1563–81, 1966.
- [11] E. Horvitz. Reasoning about beliefs and actions under computational resource constraints. In *Proceedings of the 3rd Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 301–324, 1987.
- [12] E. Horvitz. Reasoning under varying and uncertain resource constraints. In *Proceedings of the 15th National Conference on Artificial Intelligence (AAAI)*, pages 111–116, 1998.
- [13] D. G. Kirkpatrick. Hyperbolic dovetailing. In *Proceedings of the 17th Annual European Symposium on Algorithms (ESA)*, pages 616–627, 2009.

- [14] A. López-Ortiz, S. Angelopoulos, and A.M. Hamel. Optimal scheduling of contract algorithms for anytime problems. *Journal of Artificial Intelligence Research*, (51):533–554, 2014.
- [15] A. McGregor, K. Onak, and R. Panigrahy. The oil searching problem. In *Proc. of the 17th European Symposium on Algorithms (ESA)*, pages 504–515, 2009.
- [16] S. J. Russell and S. Zilberstein. Composing real-time systems. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 212–217, 1991.
- [17] S. Schuierer. Lower bounds in online geometric searching. *Computational Geometry: Theory and Applications*, 18(1):37–53, 2001.
- [18] S. Zilberstein. Using anytime algorithms in intelligent systems. *AI Magazine*, 17(3):73–83, 1996.
- [19] S. Zilberstein, F. Charpillet, and P. Chassaing. Optimal sequencing of contract algorithms. *Ann. Math. Artif. Intell.*, 39(1-2):1–18, 2003.