



HAL
open science

Contract Scheduling with Distributional and Multiple Advice

Spyros Angelopoulos, Marcin Bienkowski, Christoph Dürr, Bertrand Simon

► **To cite this version:**

Spyros Angelopoulos, Marcin Bienkowski, Christoph Dürr, Bertrand Simon. Contract Scheduling with Distributional and Multiple Advice. The 33rd International Joint Conference in Artificial Intelligence (IJCAI-24), Aug 2024, Jeju Island, South Korea. hal-04550239

HAL Id: hal-04550239

<https://hal.science/hal-04550239v1>

Submitted on 17 Apr 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.

Contract Scheduling with Distributional and Multiple Advice

Spyros Angelopoulos¹, Marcin Bienkowski², Christoph Dürr¹, Bertrand Simon³

¹Sorbonne Université, CNRS, LIP6

²University of Wrocław

³CNRS, IN2P3 Computing Center

{spyros.angelopoulos, christoph.durr}@lip6.fr, marcin.bienkowski@cs.uni.wroc.pl,
bertrand.simon@cnrs.fr

Abstract

Contract scheduling is a widely studied framework for designing real-time systems with interruptible capabilities. Previous work has showed that a prediction on the interruption time can help improve the performance of contract-based systems, however it has relied on a single prediction that is provided by a deterministic oracle. In this work, we introduce and study more general and realistic learning-augmented settings in which the prediction is in the form of a probability distribution, or it is given as a set of multiple possible interruption times. For both prediction settings, we design and analyze schedules which perform optimally if the prediction is accurate, while simultaneously guaranteeing the best worst-case performance if the prediction is adversarial. We also provide evidence that the resulting system is robust to prediction errors in the distributional setting. Last, we present an experimental evaluation that confirms the theoretical findings, and illustrates the performance improvements that can be attained in practice.

1 Introduction

A central requirement in the design of real-time and intelligent systems is the provision for anytime capabilities. Many applications, such as motion-planning and medical diagnosis, require systems that are able to output a reasonably efficient solution even if they are interrupted at arbitrary points in time. This motivates the design and evaluation of anytime systems given, as building component, more rudimentary systems that are not interruptible. Questions related to trade-offs between resources (e.g., computational time) and performance are at the heart of AI applications, and the topic of early influential works on flexible computation, resource-bounded algorithms and time-dependent planning [Boddy and Dean, 1994; Horvitz, 1988; Zilberstein and Russell, 1996].

We are interested, in particular, in a paradigm introduced by [Russell and Zilberstein, 1991], in which the building component is a *contract* algorithm, namely an anytime algorithm that is given the exact amount of allowable computation time as part of its input. A contract algorithm will always output the correct result if it is allowed at least its pre-

scribed computation time (hence its name), otherwise it may very well return a meaningless result if it is queried prior to its promised contract time. Thus, contract algorithms are not interruptible; however [Russell and Zilberstein, 1991] proposed a methodology for obtaining interruptible systems via *scheduling* consecutive executions of the contract algorithm with increasing computation times. As an example, consider a schedule in which the i -th execution of the contract algorithm is allowed time 2^i , for all $i \in \mathbb{N}$. This yields a system in which, at any time t , the contract algorithm has completed at least one execution of time $t/4$. The factor 4 quantifies the performance of the schedule, and is the multiplicative loss due to the repeated executions of the contract component.

More generally, given a contract algorithm A , a *contract schedule* is defined as an increasing sequence $X = (x_i)_{i \geq 0}$, where x_i is the length of the i -th execution of A . To evaluate the performance of a schedule X , we rely on a worst-case measure known as the *acceleration ratio* [Russell and Zilberstein, 1991]. Let $\ell(X, T)$ denote the largest contract length completed by time T in X , then the acceleration ratio of X is defined formally as

$$\text{acc}(X) = \sup_T \frac{T}{\ell(X, T)}. \quad (1)$$

It is easy to show that the optimal acceleration ratio is equal to 4, and that it is obtained by the doubling schedule $x_i = 2^i$. However, contract scheduling becomes far more challenging in more complex settings that have been studied in the literature. This includes schedules that are run on multiple processors [Bernstein *et al.*, 2002], on multiple instances [Zilberstein *et al.*, 2003], or combinations of these settings [Bernstein *et al.*, 2003; López-Ortiz *et al.*, 2014]; soft interruptions where the query time is not a hard deadline [Angelopoulos *et al.*, 2008]; performance measures beyond the acceleration ratio [Angelopoulos and López-Ortiz, 2009]; adaptive schedules [Angelopoulos and Panagiotou, 2023]; schedules with end-guarantees on completion time [Angelopoulos and Jin, 2019]; and learning-augmented schedules [Angelopoulos and Kamali, 2023]. Furthermore, contract scheduling is an abstraction of resource allocation under uncertainty, hence it has connections to other optimization problems under uncertainty, such as searching for a hidden target in a known environment under the competitive ratio, as shown in [Bernstein *et al.*, 2003; Angelopoulos, 2015].

1.1 Contract scheduling with predictions

The standard formulation of contract scheduling assumes that the scheduler has no prior information on the interruption time, hence the acceleration ratio evaluates the system performance at worst-case interruptions. In practice, however, one expects that the scheduler should be able to benefit from a *prediction* on the interruption, that may be available via a ML oracle. This motivated the recent study [Angelopoulos and Kamali, 2023] of the problem, in a setting in which an imperfect oracle provides the system with a single, deterministic prediction (advice) τ concerning the interruption time T . [Angelopoulos and Kamali, 2023] showed that there exists a schedule with acceleration ratio equal to 4 even if the prediction is adversarial, and which also yields a much improved ratio equal to 2 if the prediction is correct (i.e., error-free). Moreover, they showed that this result is optimal. Using the terminology of *learning-augmented* online algorithms [Lykouris and Vassilvitskii, 2021; Purohit *et al.*, 2018], we say that there exist schedules of *consistency* equal to 2, and *robustness* equal to 4, whereas no 4-robust schedule can be better than 2-consistent¹. We emphasize that the prediction is assumed to be *single*, i.e., it consists of a unique predicted interruption time, and deterministic, in that the oracle does not incorporate any stochastic aspects.

In practice, however, the above model may not adequately capture the information content of the anticipated prediction. For instance, in motion-planning algorithms, the system may operate under a certain probabilistic knowledge of the terrain, hence specific actions may have to be triggered according to a (stochastic) belief about the environment [Zilberstein and Russell, 1993]. For a different example, a medical diagnostic system may have to be queried at multiple anticipated times (e.g., depending on the availability of various facilities and specialists). Note that such settings are not captured by the model of [Angelopoulos and Kamali, 2023], in which the schedule is fine-tuned according to a single prediction, and may have consistency as high as 4 if the prediction includes as few as two possible interruption times.

1.2 Contribution

Motivated by the above applications, and the limitations of single/deterministic prediction oracles, we study contract scheduling under more general models that incorporate *distributional* and *multiple* advice. As in [Im *et al.*, 2023a], which studied the learning-augmented dynamic acknowledgment problem, we aim to *simultaneously* optimize the consistency and the robustness of the system. For both settings, we show that the results we obtain are tight.

We begin with the distributional setting, in which the advice oracle provides the scheduler with a distribution on the anticipated prediction. Here, the consistency is evaluated, in expectation, relative to the distributional advice (see Section 2 for the formal definition). We show how to construct a collection of n schedules for any given $n \in \mathbb{N}^*$, such that each schedule in the collection is 4-robust, and the best schedule

¹[Angelopoulos and Kamali, 2023] also studied a query-based setting in which the prediction is elicited via responses to binary queries, however, this query model is not relevant to our work.

has consistency at most $4n(2^{\frac{1}{n}} - 1)$. In particular, we show that as n grows, the consistency of the system is arbitrarily close to $4 \ln 2 \approx 2.77$, and that the best schedule can be computed in time polynomial in n . We show that this bound is optimal, in that there exists a distributional prediction for which the consistency of any 4-robust schedule is at least $4 \ln 2$. Furthermore, we demonstrate an interesting disconnect between deterministic and distributional predictions. Namely, we prove that, given the distributional prediction that maximizes the consistency, the performance of the optimal schedule deteriorates *smoothly* as function of the prediction error, measured by the Earth Mover’s Distance, or EMD [Rubner *et al.*, 1998]. In contrast, no consistency-optimal 4-robust schedule can exhibit smoothness against deterministic predictions. This disconnect shows that distributional predictions can help mitigate pathological situations, which can be of interest in other learning-augmented optimization problems.

In the second part of this work, we study the model in which the advice oracle provides the scheduler a set P of k potential interruption times (e.g., provided by k experts). Here, the consistency is measured as the worst-case performance ratio among interruptions in P , and we refer to Section 2 for the formal definition. We show how to derive a 4-robust schedule of optimal consistency $2^{2-\frac{1}{k}}$ in time $O(k \log k)$. We conclude with an experimental evaluation of our schedules, in both the distributional and multiple advice settings, that demonstrates the performance improvements that can be attained in practice.

1.3 Other related work

Motivated by the capacity of ML predictions to improve algorithmic performance, the field of learning-augmented algorithms has been growing rapidly in the recent years. We refer to the survey [Mitzenmacher and Vassilvitskii, 2020] and the online repository [Lindermayr and Megow, 2023] that lists several works over the last five years. The vast majority of works have focused on single, deterministic predictions. Multi-prediction oracles were first studied in the context of ski rental [Gollapudi and Panigrahi, 2019], followed by works on multi-shop ski rental [Wang *et al.*, 2020], facility location [Almanza *et al.*, 2021], matching and scheduling [Dinitz *et al.*, 2022], online covering [Anand *et al.*, 2022] and k -server [Antoniadis *et al.*, 2023]. Distributional predictions were first studied in [Diakonikolas *et al.*, 2021] in problems such as ski rental and prophet inequalities.

Contract scheduling is related to the *online bidding* problem [Chrobak and Kenyon-Mathieu, 2006], which is used a building block in many online algorithms, e.g., for scheduling on related machines, both non-preemptive [Aspnes *et al.*, 1997], and preemptive [Ebenlendr and Sgall, 2009], k -median problem [Chrobak *et al.*, 2008], or multi-level aggregation [Bienkowski *et al.*, 2021]. A reduction between the contract scheduling and online bidding problems for the single prediction setting was shown in [Angelopoulos and Kamali, 2023]. Consistency/robustness tradeoffs for online bidding with a single prediction were shown in [Anand *et al.*, 2021; Angelopoulos *et al.*, 2020; Im *et al.*, 2023b].

2 Preliminaries

A contract schedule is defined as a sequence of the form $X = (x_i)_{i \in \mathbb{N}}$, where x_i is the *length* of the i -th contract, and recall that the acceleration ratio of X is given by (1). In a learning-augmented setting, the acceleration ratio of X is equivalently called the *robustness* of X , and we say that X is r -robust if it has robustness at most r . From [Russell and Zilberstein, 1991] we thus know that $r \geq 4$, and that for $X = (2^i)_{i \in \mathbb{N}}$, we have that $r(X) = 4$.

The *consistency* of X is defined according to the specifics of the prediction oracle, hence we make a distinction between the two learning-augmented settings we study. In the distributional setting, the advice consists of a distribution μ on the anticipated interruption time, and the consistency of a schedule X with advice μ is defined as

$$c(X, \mu) = \frac{\mathbb{E}_{z \sim \mu}[z]}{\mathbb{E}_{z \sim \mu}[\ell(X, z)]}, \quad (2)$$

and recall that the random variable $\ell(X, z)$ denotes the largest contract completed in X by time z . We will call $\ell(X, z)$ the *profit* of X , where z is drawn from distribution μ .

In the multiple-advice setting, the prediction consists of a set $P = \{\tau_1, \dots, \tau_k\}$ of k possible interruption times (e.g., provided by k experts). Here, we measure the consistency $c(X, P)$ of a schedule X with prediction set P as its worst-case performance relative to the prediction in P , i.e., we define

$$c(X, P) = \sup_{\tau \in P} \frac{\tau}{\ell(X, \tau)}. \quad (3)$$

Note that without any assumptions, no schedule can have bounded robustness, if the interruption time is allowed to be arbitrarily small. There are two types of assumptions that can be applied to circumvent this technical issue. The first is to assume that the interruption can only occur after the first contract has completed its execution. The second is to assume that the schedule is *bi-infinite*, in that it starts with an infinite number of infinitesimally small contracts. For instance, the doubling schedule can be described as $(2^i)_{i \in \mathbb{Z}}$, and the completion time of contract $i \geq 0$ is defined as $\sum_{j=-\infty}^i 2^j = 2^{i+1}$. We choose the second assumption since it simplifies the calculations, but we note that the two assumptions can be used interchangeably; see, e.g., the discussion in [Demaine *et al.*, 2006].

For a given $\lambda \in [0, 1)$, define the schedule $X(\lambda) \triangleq (2^{i-\lambda})_{i \in \mathbb{Z}}$. The following proposition shows that it suffices to focus on the set of schedules $\cup_{\lambda \in [0, 1)} \{X(\lambda)\}$.

Proposition 1 (Appendix). *For any $\lambda \in [0, 1)$, $X(\lambda)$ is 4-robust. Conversely, every 4-robust schedule must belong in the class $\cup_{\lambda \in [0, 1)} \{X(\lambda)\}$.*

3 Distributional advice

We begin with the setting in which the advice is in the form of a given distribution μ . We first define an appropriate collection S_n of n schedules which will be instrumental towards finding an optimal schedule.

Definition 2. *For any n , let S_n denote the following collection of n schedules X_0, \dots, X_{n-1} , defined as $X_j = X(j/n) = (2^{i-j/n})_{i \in \mathbb{Z}}$.*




Figure 1: Illustration of the computation of $\ell(X_j, z)$ for $n = 3$ and $j = 1$. Note that the time scale is logarithmic. Fix an interruption point $z \in [2^2, 2^3)$. Then z is contained in I_2^k for some value of $k \in \{0, \dots, n-1\}$. If $k \geq n-j = 2$, then the largest completed contract is c_2 (of length $2^{2-1/3}$), and otherwise, the largest completed contract is c_1 (of length $2^{1-1/3}$).

Theorem 3. *For any $n \in \mathbb{N}^+$, there exists a 4-robust schedule in S_n that has consistency at most $4n \cdot (2^{1/n} - 1)$.*

Proof. First, observe that any schedule X_j satisfies the conditions of Proposition 1, and hence is 4-robust. Thus, it remains to show that for any distribution μ over $(0, +\infty)$, there exists a schedule in S_n with the desirable consistency.

We define the intervals I_i^k for $i \in \mathbb{Z}$ and $k \in \{0, \dots, n-1\}$ as $I_i^k = [2^{i+k/n}, 2^{i+(k+1)/n})$. Note that the intervals I_i^k are disjoint and their union is equal to the support of μ , i.e., $\bigcup_{i \in \mathbb{Z}} \bigcup_{0 \leq k \leq n-1} I_i^k = (0, +\infty)$.

If the interruption time z is drawn from distribution μ , then its expected value is

$$\begin{aligned} \mathbb{E}_{z \sim \mu}[z] &\leq \sum_{i \in \mathbb{Z}} \sum_{k=0}^{n-1} 2^{i+(k+1)/n} \cdot \mu(I_i^k) \\ &\leq \sum_{i \in \mathbb{Z}} 2^i \cdot \sum_{k=0}^{n-1} \mu(I_i^k) \cdot 2^{k/n} \cdot 2^{1/n}. \end{aligned}$$

We now estimate the expected value of $\ell(X_j, z)$ for each schedule X_j with $j \in \{0, \dots, n-1\}$. To this end, we consider a fixed value of interruption time z and let I_i^k be the interval containing z . We note that if $k \geq n-j$, then $2^{i+1-j/n} \leq z < 2^{i+2-j/n}$, and thus $\ell(X_j, z) = 2^{i-j/n}$. Otherwise, i.e., if $k < n-j$, we have $\ell(X_j, z) = 2^{i-1-j/n}$. See Figure 1 for a pictorial description of this computation.

Hence,

$$\begin{aligned} \mathbb{E}_{z \sim \mu}[\ell(X_j, z)] &= \sum_{i \in \mathbb{Z}} \sum_{k=0}^{n-j-1} 2^{i-1-j/n} \cdot \mu(I_i^k) \\ &\quad + \sum_{i \in \mathbb{Z}} \sum_{k=n-j}^{n-1} 2^{i-j/n} \cdot \mu(I_i^k). \end{aligned}$$

Next, we compute the sum of the largest contracts over all schedules X_j , as follows:

$$\begin{aligned} &\sum_{j=0}^{n-1} \mathbb{E}_{z \sim \mu}[\ell(X_j, z)] \\ &= \sum_{j=0}^{n-1} \sum_{i \in \mathbb{Z}} \left(\sum_{k=0}^{n-j-1} 2^{i-1-\frac{j}{n}} \cdot \mu(I_i^k) + \sum_{k=n-j}^{n-1} 2^{i-\frac{j}{n}} \cdot \mu(I_i^k) \right) \\ &= \sum_{i \in \mathbb{Z}} 2^i \cdot \sum_{k=0}^{n-1} \mu(I_i^k) \cdot \left(\frac{1}{2} \sum_{j=0}^{n-k-1} 2^{-j/n} + \sum_{j=n-k}^{n-1} 2^{-j/n} \right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{i \in \mathbb{Z}} 2^i \cdot \sum_{k=0}^{n-1} \mu(I_i^k) \cdot \left(\frac{1}{2} \cdot \frac{1 - 2^{\frac{k-n}{n}}}{1 - 2^{-\frac{1}{n}}} + 2^{\frac{k-n}{n}} \cdot \frac{1 - 2^{-\frac{k}{n}}}{1 - 2^{-\frac{1}{n}}} \right) \\
&= \sum_{i \in \mathbb{Z}} 2^i \cdot \sum_{k=0}^{n-1} \mu(I_i^k) \cdot \frac{2^{k/n-2}}{1 - 2^{-1/n}} \\
&= \frac{\mathbb{E}_{z \sim \mu}[z]}{2^{1/n}} \cdot \frac{2^{-2}}{1 - 2^{-1/n}} = \frac{\mathbb{E}_{z \sim \mu}[z]}{4(2^{1/n} - 1)}.
\end{aligned}$$

We can now use an averaging argument, and deduce that there exists a schedule X_{j^*} in S_n such that

$$\begin{aligned}
\mathbb{E}_{z \sim \mu}[\ell(X_{j^*}, z)] &\geq \frac{1}{n} \cdot \sum_{j=0}^{n-1} \mathbb{E}_{z \sim \mu}[\ell(X_j, z)] \\
&\geq \frac{\mathbb{E}_{z \sim \mu}[z]}{4n(2^{1/n} - 1)}.
\end{aligned}$$

Hence, the consistency of the schedule X_{j^*} is $c(X_{j^*}, \mu) = \mathbb{E}_{z \sim \mu}[z] / \mathbb{E}_{z \sim \mu}[\ell(X_{j^*}, z)] \leq 4n(2^{1/n} - 1)$. \square

We observe that $4n(2^{1/n} - 1)$ is decreasing and tends to $4 \ln 2$ as $n \rightarrow \infty$. Indeed, $4n(2^{1/n} - 1) = 4n(e^{\ln 2 \cdot (1/n)} - 1) = 4n(\ln 2 \cdot (1/n) + O(1/n^2)) = 4 \ln 2 + O(1/n)$. This means that for any $\epsilon > 0$, there exists an integer $n = O(1/\epsilon)$ such that $4n(2^{1/n} - 1) \leq 4 \ln 2 + \epsilon$. Assuming that we can evaluate the probability associated with an interval in constant time, we obtain the following corollary.

Corollary 4. *For any arbitrarily small $\epsilon > 0$, there is an algorithm with runtime polynomial in $O(1/\epsilon)$ for devising a 4-robust schedule that has consistency at most $4 \cdot \ln 2 + \epsilon$.*

We now show that the upper bound of Theorem 3 is tight.

Theorem 5. *For any $n \in \mathbb{N}$, there exists a distributional prediction for which every collection C_n that consists of n 4-robust schedules cannot contain a schedule of consistency smaller than $4n \cdot (2^{1/n} - 1)$.*

Proof. Using Proposition 1, let us denote, without loss of generality, the 4-robust schedules in C_n by $X(\lambda_k) = (2^{i-\lambda_k})_{i=-\infty}^{\infty}$ for $k \in [0, n-1]$ and $0 \leq \lambda_k < \lambda_{k+1} < 1$.

We will define a distribution μ that has an n -point support, and where the k -th point occurs at time $2^{2-\lambda_k-\epsilon}$ with probability p_k , for $\epsilon > 0$ arbitrarily small. For $k < n-1$, we define $p_k = 2^{\lambda_{k+1}-\lambda_0} - 2^{\lambda_k-\lambda_0}$, and define $p_{n-1} = 2 - 2^{\lambda_{n-1}-\lambda_0}$. Let $P_k = \sum_{j \leq k} p_j$, so $P_k = 2^{\lambda_{k+1}-\lambda_0} - 1$.

We have that $\mathbb{E}_{z \sim \mu}[\ell(X(\lambda_0), z)] = 2^{-\lambda_0}$ as the support of μ is in the interval $[2^{1-\lambda_0}, 2^{2-\lambda_0}]$. Therefore, if an interruption occurs at any point in the support of μ , the last completed contract of $X(\lambda_0)$ is of length $2^{-\lambda_0}$.

Similarly, for any $k > 0$, we have:

$$\begin{aligned}
\mathbb{E}_{z \sim \mu}[\ell(X_{\lambda_k}, z)] &= P_{k-1} 2^{1-\lambda_k} + (1 - P_{k-1}) 2^{-\lambda_k} \\
&= 2^{1-\lambda_0} - 2^{1-\lambda_k} + 2^{1-\lambda_k} - 2^{-\lambda_0} = 2^{-\lambda_0}. \quad (4)
\end{aligned}$$

The expected value of μ is (defining $\lambda_n = \lambda_0 + 1$):

$$\begin{aligned}
\mathbb{E}_{z \sim \mu}[z] &= \sum_{j < n} (2^{\lambda_{j+1}-\lambda_0} - 2^{\lambda_j-\lambda_0}) \cdot 2^{2-\lambda_j} \\
&= 2^{-\lambda_0} \cdot \sum_{j < n} (2^{2+\lambda_{j+1}-\lambda_j} - 2^2) \\
&= 2^{2-\lambda_0} \cdot \left(-n + \sum_{j < n} 2^{\lambda_{j+1}-\lambda_j} \right). \quad (5)
\end{aligned}$$

From (4) and (5), the consistency of any $X(\lambda_k)$ is:

$$c(X(\lambda_k), \mu) = 4 \cdot \left(-n + \sum_{j < n} 2^{\lambda_{j+1}-\lambda_j} \right).$$

$c(X(\lambda_k), \mu)$ is minimized when the λ_j are chosen so as to minimize the sum term, while respecting the constraint that the sum of all the exponents equals $\lambda_n - \lambda_0 = 1$ and each exponent belongs to $[0, 1)$. By the convexity of the function $x \mapsto 2^x$, this occurs when all terms are equal to $2^{1/n}$, hence

$$c(X(\lambda_k), \mu) \geq 4 \cdot \left(-n + n \cdot 2^{1/n} \right) = 4n \cdot (2^{1/n} - 1). \quad \square$$

By allowing $n \rightarrow \infty$, Theorem 5 shows that no 4-robust schedule can have consistency better than $4 \ln 2$. Note that the theorem relies on a probability distribution with n mass points. However, we can obtain the same lower bound by relying to a prediction given as a continuous distribution.

Theorem 6 (Appendix). *For any $D > 0$, there is a continuous distribution μ_D over the interval $[D, 2D]$ such that, for every 4-robust schedule X , $c(X, \mu_D) \geq 4 \ln 2$.*

3.1 Smoothness with prediction errors

We will now discuss an interesting disconnect between deterministic and distributional predictions that arises in contract scheduling. Consider first the setting of a single, deterministic prediction, for which we know that there exist 4-robust, 2-consistent schedules [Angelopoulos and Kamali, 2023]. Given such prediction, say τ , let $\eta = |T - \tau|$ denote the prediction error associated with τ , where T denotes the actual interruption.

Proposition 7 (Appendix). *For any arbitrarily small $\epsilon > 0$, and any schedule X with prediction τ that is 4-robust and 2-consistent, there exists T such that $\eta = |T - \tau| = \epsilon$ and $\ell(X, T) \leq \frac{1}{4}T + \frac{\epsilon}{4}$.*

Proposition 7 shows that, in the single prediction setting, any schedule that simultaneously optimizes the consistency and the robustness is extremely fragile with respect to prediction errors. Informally, in the presence of a tiny prediction error, the proposition shows that the consistency of the schedule becomes as large as the robustness, namely arbitrarily close to 4, hence the single prediction leads to no improvement in any practical situation.

In contrast, we will show that this is not the case for the worst-case distributional prediction of Theorem 6. Specifically, in Theorem 8, we will show that the performance of the schedule degrades smoothly as a function of the prediction error. For some intuition behind the proof, we will use the fact that worst-case distributional predictions require that

a large time interval contributes to the schedule's expected profit, and that the distribution is rather "balanced" over that interval. Therefore, if an adversary were to change substantially the performance of the schedule, by altering the predicted distribution, it would have to shift a large amount of probability mass over a long time span. This motivates the choice of the well-known Earth Mover's Distance (EMD) as a metric of the prediction error. We will show that a small EMD error has a likewise small effect on the performance.

Given two probability distributions μ and μ' over \mathbb{R}^+ , the Earth Mover's Distance (EMD) [Rubner *et al.*, 1998] intuitively represents the minimum cost incurred in order to obtain μ' from μ , where moving an infinitesimal probability mass costs its amount multiplied by the distance moved. Formally, let $\Psi(\mu, \mu') \subset (\mathbb{R}^+ \times \mathbb{R}^+) \rightarrow \mathbb{R}^+$ be the set of functions ψ such that $\forall x \in \mathbb{R}^+, \int_{\mathbb{R}^+} \psi(x, y) dy = \mu(x)$ and $\forall y \in \mathbb{R}^+, \int_{\mathbb{R}^+} \psi(x, y) dx = \mu'(y)$. Here, $\psi(x, y)$ describes the amount of mass that is moved from point x in μ to point y in μ' . The distance $\text{EMD}(\mu, \mu')$ is then defined as

$$\text{EMD}(\mu, \mu') = \inf_{\psi \in \Psi(\mu, \mu')} \int_{\mathbb{R}^+} \int_{\mathbb{R}^+} |x - y| \psi(x, y) dx dy.$$

Theorem 8. *Given $D \in \mathbb{R}^+$, let μ_D be the probability distribution of Theorem 6. Then, for any probability distribution μ' such that $\text{EMD}(\mu_D, \mu') \leq \eta$, for sufficiently small η , any 4-robust schedule X with prediction μ_D satisfies*

$$\mathbb{E}_{z \sim \mu'}[\ell(X, z)] \geq \frac{\mathbb{E}_{z \sim \mu'}[z]}{4 \ln 2 + O(\sqrt{\eta/D})}.$$

Proof Sketch. We give the main elements of the proof, and defer some technical parts to the appendix.

Recall that the distribution $\mu = \mu_D$ is defined over $[D, 2D]$ by the density function $x \mapsto 2D/x^2$. This density function takes values in the interval $[1/(2D), 2/D]$. From Proposition 1, we know that any 4-robust schedule X is of the form $X(\lambda) = 2^{i-\lambda}$, for some $\lambda \in [0, 1)$. For given $\eta > 0$, let μ' be any probability distribution such that $\text{EMD}(\mu, \mu') \leq \eta$.

By the definition of EMD, let $\psi \in \Psi(\mu, \mu')$ be a function such that

$$\text{EMD}(\mu, \mu') = \int_{\mathbb{R}^+} \int_{\mathbb{R}^+} |x - y| \psi(x, y) dx dy.$$

We decompose ψ in two functions ψ^+ and ψ^- satisfying the following conditions. For $x < y$, $\psi^+(x, y) = \psi(x, y)$ and $\psi^-(x, y) = 0$. For $x > y$, $\psi^-(x, y) = \psi(x, y)$ and $\psi^+(x, y) = 0$. Last, we require that $\psi^+(x, x) = \int_{z \leq x} \psi(x, z) dz$, and $\psi^-(x, x) = \int_{z \geq x} \psi(x, z) dz$. Intuitively, ψ^+ denotes the probability mass moved towards higher values, and, likewise, ψ^- denotes the probability mass moved towards lower values. Let μ^+ be the probability obtained after applying ψ^+ ; namely, we define $\mu^+(y) = \int_{\mathbb{R}^+} \psi^+(x, y) dx$. Note that

$$\eta \geq \text{EMD}(\mu, \mu') = \text{EMD}(\mu, \mu^+) + \text{EMD}(\mu^+, \mu').$$

In the first part of the proof (details in the appendix), we focus on the effect of ψ^+ and bound the expected profit of $X(\lambda)$ under the distribution μ^+ , as well as the expectation of

the distribution μ^+ . The intuition here is that the expected profit of $X(\lambda)$ cannot decrease, and that the growth of the expected value of the distribution is bounded by the EMD. We show

$$\mathbb{E}_{z \sim \mu^+}[\ell(X(\lambda), z)] \geq \mathbb{E}_{z \sim \mu}[\ell(X(\lambda), z)], \quad (6)$$

$$\mathbb{E}_{z \sim \mu^+}[z] \leq \mathbb{E}_{z \sim \mu}[z] + \eta. \quad (7)$$

In the next part of the proof (details in the appendix), we focus on the effect of ψ^- . Here, the expected value of the distribution cannot increase, and the expected profit of $X(\lambda)$ decreases by an amount which can be bounded using η and D . This last bound is the most technical step of the proof and uses the fact that $\min_{x \in [D, 2D]} \mu(x) \leq 4 \cdot \max_{x \in [D, 2D]} \mu(x)$. We show

$$\mathbb{E}_{z \sim \mu'}[z] \leq \mathbb{E}_{z \sim \mu^+}[z], \quad (8)$$

$$\mathbb{E}_{z \sim \mu'}[\ell(X(\lambda), z)] \geq \mathbb{E}_{z \sim \mu^+}[\ell(X(\lambda), z)] - 8\sqrt{2\eta D}. \quad (9)$$

To conclude the proof, note that $\mathbb{E}_{z \sim \mu}[z] = 2D \ln 2$ and $\mathbb{E}_{z \sim \mu}[\ell(X(\lambda), z)] = D/2$, from the proof of Theorem 6. Combining Equations (6), (7), (8) and (9), we arrive at the desired result, for $\eta < D/512$ (to ensure there is no division by zero). See Appendix for details. \square

4 Multiple advice

In this section, we study the setting in which the prediction oracle provides a set P that consists of k (potential interruption) times, denoted by $\tau_0, \dots, \tau_{k-1}$. Recall that the consistency of a schedule is given by (3). For every $j \in [0, k-1]$, we denote by $\delta_j \in [0, 1)$, and by $i_j \in \mathbb{N}$ the unique values such that $\tau_i = 2^{i_j + \delta_j}$. We assume, without loss of generality that $0 \leq \delta_0 \leq \delta_1 \dots \leq \delta_{k-1} < 1$.

We first show that we can find, in time polynomial in k , a schedule that simultaneously minimizes the consistency and the robustness.

Theorem 9. *Given a prediction set P of size k , we can find a 4-robust schedule of optimal consistency in time $O(k^2)$.*

Proof. We first argue that there exists a 4-robust schedule of optimal consistency $X^* = (x_i^*)_{i \in \mathbb{Z}}$, such that X^* contains at least one contract whose execution is completed at time τ_j , for some $\tau_j \in P$. By way of contradiction, suppose this is not the case. From Proposition 1, the last contract completed by X^* by time t finishes at time $2\ell(X^*, t)$. Define $a = \min_{j \in [0, k-1]} \frac{\tau_j}{2\ell(X^*, \tau_j)}$, and consider the schedule $X' = (2^{\log_2 a} \cdot x_i^*)_{i \in \mathbb{Z}}$. From Proposition 1, X' is also 4-robust. Moreover, we have that $\ell(X', \tau_j) \geq \ell(X^*, \tau_j)$, therefore X' is also optimal, a contradiction.

The above observation leads to the following algorithm for finding an optimal schedule. For every $j \in [0, k-1]$, define the schedule $X_j = (2^{i+\delta_j})_{i \in \mathbb{Z}}$, and note that, by definition, X_j has a contract that terminates at time τ_j . Thus, among schedules in the collection $\{X_j\}_{j=0}^{k-1}$, the schedule X_{l^*} with best consistency is such that

$$l^* = \arg \min_{j \in [0, k-1]} \alpha_j, \text{ where } \alpha_j = \max_{i \in [0, k-1]} \frac{\tau_i}{\ell(X_j, \tau_i)}.$$




Figure 2: Illustration of the definitions of δ_j and D_j .

From Proposition 1, it follows that the above algorithm yields a 4-robust schedule of optimal consistency. \square

The complexity of this algorithm is $O(k^2)$, since each α_j can be computed in time $O(k)$. However, we can reduce the complexity to $O(k \log k)$, using an argument that will also be useful in the proof of Corollary 11, and which is illustrated in Figure 2. Consider the circle of unit perimeter, with an arbitrary point O fixed, and let p_0, \dots, p_{k-1} be points in the circle, such that the clockwise (arc-length) distance between O and p_j is $d_c(O, p_j) = \delta_j$. Given $i \in [0, k-1]$, define D_i as the clockwise distance between the consecutive points $p_{(i-1) \bmod k}$ and p_i in the circle. We show the following:

Lemma 10. *For $j \in [0, k-1]$, let X_j denote the schedule $(2^{i+\delta_j})_{i \in \mathbb{Z}}$. Then X_j has consistency at most 2^{2-D_j} .*

Proof. From construction, it readily follows that the maximum consistency $c(X_j, P)$ is attained when the interruption occurs at time $\tau_{(j-1) \bmod k} \in P$. We consider two cases. If $j \neq 0$, then $(j-1) \bmod k = j-1$, and by definition, we have that $\ell(X_j, \tau_{j-1}) = 2^{i_{j-1}-2+\delta_j}$. Therefore,

$$c(X_j, P) \leq \frac{2^{i_{j-1}+\delta_{j-1}}}{2^{i_{j-1}-2+\delta_j}} = 2^{2-D_j}.$$

For the second case, suppose that $j = 0$, thus $(j-1) \bmod k = k-1$. Then, $\ell(X_0, \tau_{k-1}) = 2^{i_{k-1}+\delta_0-1}$. Therefore,

$$c(X_0, P) \leq \frac{2^{i_{k-1}+\delta_{k-1}}}{2^{i_{k-1}+\delta_0-1}} = 2^{1+(\delta_{k-1}-\delta_0)} = 2^{2-D_0},$$

which concludes the proof. \square

Lemma 10 implies that we can improve the complexity of the algorithm of Theorem 9, by finding the index j for which D_j is minimized. This can be accomplished in time $O(k \log k)$ by sorting.

An interesting question is finding the exact value of the worst-case consistency of optimal schedules. The following corollary answers this question.

Corollary 11. *The schedule of Theorem 9 has consistency at most $2^{2-\frac{1}{k}}$, where k is the size of P . Furthermore, this bound is tight, in that there exists a prediction P such that every 4-robust schedule has consistency at least $2^{2-\frac{1}{k}}$.*

Proof. From Lemma 10, since the smallest clockwise distance between two consecutive points in the unit circle cannot exceed $1/k$, we obtain that the worst-case consistency is at most $2^{2-\frac{1}{k}}$. This is tight, if the k points are equidistantly distributed, i.e., if $\delta_j = j/k$, for every $j \in [0, k-1]$. \square

Note that Corollary 11 subsumes the known results for the extreme cases $k = 1$ (for which the consistency is equal to 2 [Angelopoulos and Kamali, 2023]), and $k \rightarrow \infty$ (for which the consistency reduces to the worst-case acceleration ratio [Russell and Zilberstein, 1991]).

5 Experimental evaluation

We present an experimental evaluation of our schedules, for both advice settings. We report the main findings, and we refer to the Appendix for additional results and discussion.

5.1 Distributional advice

We evaluate our algorithm of Theorem 3, to which we refer as SEL_n , and recall that this algorithm selects the best schedule in class S_n , for some given $n \in \mathbb{N}$. We first consider, as distributional advice μ , a *normal* distribution that is truncated at zero, with mean m , and standard deviation σ .

Figure 3 depicts the experimental performance of SEL_n , as function of m , for $n \in [1, 4]$, $\sigma = 0.05m$ and $m \in [1, 1024]$. Note that for all sufficiently large m , e.g., $m \geq 100$, the expected interruption time is extremely close to m . We observe that the consistency of SEL_n improves as n increases, by design of the algorithm. In particular, for $n = 4$, the empirical consistency is 2.51. This value is, as expected, below the anticipated worst-case bound of Theorem 3 (namely, $16(2^{0.25} - 1) \approx 3.03$) and above the lower bound of 2, which applies to single, deterministic predictions.

The shape of the consistency is a saw-like function of the mean, which is explained by the fact that the consistency (as the acceleration ratio) has transitions at “critical” times, i.e., right after the completion of a contract. As n increases, we note that the transitions become more smooth. This is because the number of candidate schedules in S_n becomes larger, hence the probability that the interruption is critical for some schedule in S_n decreases. This also explains why the number of peaks increases as function of n .

We also report results for a *uniform* distributional advice. Specifically, Figure 4 depicts the consistency of SEL_n with advice chosen according to $U[0.95t, 1.05t]$, as function of t . For $n = 4$, we observe a similar empirical consistency value as in the case of the normal distribution, namely 2.44. The shape of the plot has sharper transitions, relatively to the normal distribution, which can be explained by the fact that the latter spreads the random interruption over a larger interval than the uniform one, in this experimental setup. This also explains why the difference between the peaks and valleys is more pronounced in Figure 4 than in Figure 3.

5.2 Multiple advice

We evaluate our algorithm from Section 4, to which we refer as MULT_k , where k is the cardinality of the prediction set P . Specifically, for every $k \in [1, 10]$, we generate P as k




Figure 3: Plot of the consistency of SEL_n with advice a truncated normal distribution as function of the mean m .




Figure 4: Plot of the consistency of SEL_n with advice a uniform distribution in $[0.95t, 1.05t]$, as function of t .




Figure 5: Experimental evaluation of MULT_k .

values chosen independently and uniformly at random in the interval $[1, 1024]$. For each such prediction set, we compute the worst-case consistency, as evaluated by (3). We compute, in addition, the average-case consistency, namely the average of the ratios $\tau/\ell(\text{MULT}_k, \tau)$, where τ is chosen uniformly at random from P . The latter is a much more relaxed performance measure that treats each of the k predictions in P as equally likely. We repeat this experiment 1000 times, and we report the average of the corresponding ratios.

We observe that for all values of k , both the worst-case and the average consistency are below the upper bound of $2^{2-\frac{1}{k}}$, which confirms the result of Corollary 11. The ratios, as expected, are increasing functions of k , since the larger the size of P , the more noisy the quality of the prediction.

6 Conclusion

We studied a classic optimization problem related to bounded-resource reasoning, namely contract scheduling, in novel learning-augmented settings that capture distributional and multiple predictions. In both settings, we gave analytically optimal schedules that simultaneously optimize the consistency and the robustness. As discussed in Section 1, con-

tract scheduling has been studied in a variety of settings, including multiple instances that must be solved concurrently, in a single or multiple processors. Future work will address these more complex variants under similar learning-enhanced models. It will be also interesting to perform a multi-objective analysis, in the multiple-advice model, that addresses the average expected consistency or trade-offs between the worst-case and the average-case metrics. The techniques of Section 4 can be applicable, by showing first tradeoffs between average and worst-case distances between points on the unit circle.

Another direction for future work is searching for a hidden target, with distributional or multiple predictions about its position in the environment. This setting has applications in robotic search and exploration, and has been studied with single, or no prediction e.g., [Eberle *et al.*, 2022; Sung and Tokekar, 2019]. In particular, the techniques we developed in this work will be very useful in the context of the *line* and *star* search environments [Jaillet and Stafford, 1993], given the connections between contract scheduling and competitive search [Bernstein *et al.*, 2003; Angelopoulos, 2015]. Last, our work is the first to bring attention to the fact that single predictions may be fragile for certain problems, a finding that can have implications in other domains and applications such as learning-augmented online conversion problems [Sun *et al.*, 2021].

7 Acknowledgements

This work was partially funded by the project PREDICTIONS, grant ANR-23-CE48-0010 and the project ALGORIDAM, grant ANR-19-CE48-0016 from the French National Research Agency (ANR), as well as by Polish National Science Centre grant 2022/45/B/ST6/00559.

References

- [Almanza *et al.*, 2021] Matteo Almanza, Flavio Chierichetti, Silvio Lattanzi, Alessandro Panconesi, and Giuseppe Re. Online facility location with multiple advice. *Advances in Neural Information Processing Systems*, 34:4661–4673, 2021.
- [Alpern and Gal, 2003] Steve Alpern and Shmuel Gal. *The theory of search games and rendezvous*. Kluwer Academic Publishers, 2003.
- [Anand *et al.*, 2021] Keerti Anand, Rong Ge, Amit Kumar, and Debmalya Panigrahi. A regression approach to learning-augmented online algorithms. *Advances in Neural Information Processing Systems*, 34:30504–30517, 2021.
- [Anand *et al.*, 2022] Keerti Anand, Rong Ge, Amit Kumar, and Debmalya Panigrahi. Online algorithms with multiple predictions. In *International Conference on Machine Learning*, pages 582–598. PMLR, 2022.
- [Angelopoulos and Jin, 2019] Spyros Angelopoulos and Shendan Jin. Earliest-completion scheduling of contract algorithms with end guarantees. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 5493–5499, 2019.
- [Angelopoulos and Kamali, 2023] Spyros Angelopoulos and Shahin Kamali. Contract scheduling with predictions. *J. Artif. Intell. Res.*, 77:395–426, 2023.
- [Angelopoulos and López-Ortiz, 2009] Spyros Angelopoulos and Alejandro López-Ortiz. Interruptible algorithms for multi-problem solving. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 380–386, 2009.
- [Angelopoulos and Panagiotou, 2023] Spyros Angelopoulos and Konstantinos Panagiotou. Weighted online search. *J. Comput. Syst. Sci.*, 138:103457, 2023.
- [Angelopoulos *et al.*, 2008] Spyros Angelopoulos, Alejandro López-Ortiz, and Angele Hamel. Optimal scheduling of contract algorithms with soft deadlines. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 868–873, 2008.
- [Angelopoulos *et al.*, 2020] Spyros Angelopoulos, Christoph Dürr, Shendan Jin, Shahin Kamali, and Marc P. Renault. Online computation with untrusted advice. In *Proceedings of the 11th International Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 52:1–52:15, 2020.
- [Angelopoulos, 2015] Spyros Angelopoulos. Further connections between contract-scheduling and ray-searching problems. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1516–1522, 2015.
- [Antoniadis *et al.*, 2023] Antonios Antoniadis, Christian Cooper, Marek Elias, Adam Polak, and Bertrand Simon. Mixing predictions for online metric algorithms. In *International Conference on Machine Learning*, pages 969–983. PMLR, 2023.
- [Aspnes *et al.*, 1997] James Aspnes, Yossi Azar, Amos Fiat, Serge A. Plotkin, and Orli Waarts. On-line routing of virtual circuits with applications to load balancing and machine scheduling. *J. ACM*, 44(3):486–504, 1997.
- [Bernstein *et al.*, 2002] Daniel S. Bernstein, T. J. Perkins, Shlomo Zilberstein, and Lev Finkelstein. Scheduling contract algorithms on multiple processors. In *Proceedings of the 18th AAAI Conference on Artificial Intelligence (AAAI)*, pages 702–706, 2002.
- [Bernstein *et al.*, 2003] 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.
- [Bienkowski *et al.*, 2021] Marcin Bienkowski, Martin Böhm, Jaroslav Byrka, Marek Chrobak, Christoph Dürr, Lukáš Folwarczný, Lukasz Jez, Jirí Sgall, Kim Thang Nguyen, and Pavel Veselý. New results on multi-level aggregation. *Theor. Comput. Sci.*, 861:133–143, 2021.
- [Boddy and Dean, 1994] Mark Boddy and Thomas L. Dean. Deliberation scheduling for problem solving in time-constrained environments. *Artif. Intell.*, 67(2):245–285, 1994.
- [Chrobak and Kenyon-Mathieu, 2006] Marek Chrobak and Claire Kenyon-Mathieu. SIGACT news online algorithms column 10: Competitiveness via doubling. *SIGACT News*, 37(4):115–126, 2006.
- [Chrobak *et al.*, 2008] Marek Chrobak, Claire Kenyon, John Noga, and Neal E. Young. Incremental medians via online bidding. *Algorithmica*, 50(4):455–478, 2008.
- [Demaine *et al.*, 2006] E.D. Demaine, S.P. Fekete, and S. Gal. Online searching with turn cost. *Theoretical Computer Science*, 361:342–355, 2006.
- [Diakonikolas *et al.*, 2021] Ilias Diakonikolas, Vasilis Kontonis, Christos Tzamos, Ali Vakilian, and Nikos Zari-fis. Learning online algorithms with distributional advice. In *International Conference on Machine Learning*, pages 2687–2696. PMLR, 2021.
- [Dinitz *et al.*, 2022] Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii. Algorithms with prediction portfolios. *Advances in neural information processing systems*, 35:20273–20286, 2022.
- [Ebenlendr and Sgall, 2009] Tomás Ebenlendr and Jirí Sgall. Optimal and online preemptive scheduling on uniformly related machines. *J. Sched.*, 12(5):517–527, 2009.
- [Eberle *et al.*, 2022] Franziska Eberle, Alexander Linder-mayr, Nicole Megow, Lukas Nölke, and Jens Schlöter. Robustification of online graph exploration methods. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence, AAAI 2022*, pages 9732–9740. AAAI Press, 2022.
- [Gollapudi and Panigrahi, 2019] Sreenivas Gollapudi and Debmalya Panigrahi. Online algorithms for rent-or-buy with expert advice. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 2319–2327, 2019.

- [Horvitz, 1988] Eric Horvitz. Reasoning about beliefs and actions under computational resource constraints. *Int. J. Approx. Reasoning*, 2(3):337–338, 1988.
- [Im *et al.*, 2023a] Sungjin Im, Benjamin Moseley, Chenyang Xu, and Ruilong Zhang. Online dynamic acknowledgment with learned predictions. In *INFOCOM*, pages 1–10. IEEE, 2023.
- [Im *et al.*, 2023b] Sungjin Im, Benjamin Moseley, Chenyang Xu, and Ruilong Zhang. Online state exploration: Competitive worst case and learning-augmented algorithms. In *ECML/PKDD (4)*, volume 14172 of *Lecture Notes in Computer Science*, pages 333–348. Springer, 2023.
- [Jaillet and Stafford, 1993] P. Jaillet and M. Stafford. Online searching. *Operations Research*, 49:234–244, 1993.
- [Lindermayr and Megow, 2023] Alexander Lindermayr and Nicole Megow. Repository of works on algorithms with predictions. <https://algorithms-with-predictions.github.io>, 2023. Accessed: 2023-12-01.
- [López-Ortiz *et al.*, 2014] Alejandro López-Ortiz, Spyros Angelopoulos, and Angele Hamel. Optimal scheduling of contract algorithms for anytime problem-solving. *J. Artif. Intell. Res.*, 51:533–554, 2014.
- [Lykouris and Vassilvitskii, 2021] Thodoris Lykouris and Sergei Vassilvitskii. Competitive caching with machine learned advice. *J. ACM*, 68(4):24:1–24:25, 2021.
- [Mitzenmacher and Vassilvitskii, 2020] Michael Mitzenmacher and Sergei Vassilvitskii. Algorithms with predictions. In *Beyond the Worst-Case Analysis of Algorithms*, pages 646–662. Cambridge University Press, 2020.
- [Purohit *et al.*, 2018] Manish Purohit, Zoya Svitkina, and Ravi Kumar. Improving online algorithms via ML predictions. In *Proceedings of the 31st Annual Conference on Neural Information Processing Systems (NIPS)*, pages 9661–9670, 2018.
- [Rubner *et al.*, 1998] Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. A metric for distributions with applications to image databases. In *Sixth international conference on computer vision (IEEE Cat. No. 98CH36271)*, pages 59–66. IEEE, 1998.
- [Russell and Zilberstein, 1991] Stuart J. Russell and Shlomo Zilberstein. Composing real-time systems. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 212–217, 1991.
- [Sun *et al.*, 2021] Bo Sun, Russell Lee, Mohammad Hajiesmaili, Adam Wierman, and Danny Tsang. Pareto-optimal learning-augmented algorithms for online conversion problems. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 10339–10350, 2021.
- [Sung and Tokekar, 2019] Yoonchang Sung and Pratap Tokekar. A competitive algorithm for online multi-robot exploration of a translating plume. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3391–3397. IEEE, 2019.
- [Wang *et al.*, 2020] Shufan Wang, Jian Li, and Shiqiang Wang. Online algorithms for multi-shop ski rental with machine learned advice. *Advances in Neural Information Processing Systems*, 33:8150–8160, 2020.
- [Zilberstein and Russell, 1993] Shlomo Zilberstein and Stuart Russell. Anytime sensing planning and action: A practical model for robot control. In *IJCAI*, pages 1402–1407. Morgan Kaufmann, 1993.
- [Zilberstein and Russell, 1996] Shlomo Zilberstein and Stuart J. Russell. Optimal composition of real-time systems. *Artif. Intell.*, 82(1-2):181–213, 1996.
- [Zilberstein *et al.*, 2003] Shlomo Zilberstein, Francois Charpillat, and Philippe Chassaing. Optimal sequencing of contract algorithms. *Ann. Math. Artif. Intell.*, 39(1-2):1–18, 2003.

Appendix

A Omitted proofs and proof details

Proof of Proposition 1. For the first part of the proposition, we know from [Russell and Zilberstein, 1991] that the robustness of a schedule (i.e., its acceleration ratio) is maximized for interruptions that occur infinitesimally prior to the completion time of a contract in the schedule. Thus, an equivalent expression of the robustness of a schedule X is given by

$$r(X) = \sup_i \frac{\sum_{j=-\infty}^i x_j}{x_{i-1}}. \quad (10)$$

It is then straightforward to confirm that any schedule $X(\lambda)$ is 4-robust, since it is a scaled variant of the doubling schedule, which is 4-robust. For the second part of the proposition, the result follows from the study of the linear recurrence inequality $\sum_{j \leq i} x_j \leq 4x_{i-1}$, and in particular Theorem 9.4 in [Alpern and Gal, 2003], with $M = 2$. The result is described in terms of a search problem, but the same recurrence inequality applies, and hence the same conclusion. \square

Proof of Theorem 6. From Proposition 7, let $X(\lambda)$ denote the 4-robust schedule, for some $\lambda \in [0, 1)$. Consider any value $D > 1$. We will define a distribution function μ_D over $[D, 2D]$ such that $c(X(\lambda), \mu_D) \geq 4 \ln 2$.

Specifically, let the density function of μ_D equal $f_D(x) = 2D/x^2$ on $[D, 2D]$ and 0 elsewhere. Note that $\int f_D = 1$, so μ_D is indeed a distribution. We have

$$\mathbb{E}_{z \sim \mu_D}[z] = \int_D^{2D} \frac{2D}{x} dx = 2D \ln 2.$$

Let k be such that $2^{k-\lambda} \in [D, 2D)$. The last completed contract of $X(\lambda)$ is of length $2^{k-\lambda-2}$ if an interruption occurs between D and $2^{k-\lambda}$, and of length $2^{k-\lambda-1}$ if an interruption occurs between $2^{k-\lambda}$ and $2D$. So

$$\begin{aligned} & \mathbb{E}_{z \sim \mu_D}[\ell(X(\lambda), z)] \\ &= 2^{k-2-\lambda} \int_D^{2^{k-\lambda}} \frac{2D}{x^2} dx + 2^{k-1-\lambda} \int_{2^{k-\lambda}}^{2D} \frac{2D}{x^2} dx \\ &= 2^{k-2-\lambda} (2 - D \cdot 2^{1-k+\lambda} + 2(2D \cdot 2^{\lambda-k} - 1)) \\ &= 2^{k-2-\lambda} (D \cdot 2^{1-k+\lambda}) \\ &= D/2. \end{aligned}$$

Therefore, we have $c(X(\lambda), \mu_D) = 4 \ln 2$. \square

Proof of Proposition 7. Let X denote a 4-robust, 2-consistent schedule X . Then for any given prediction τ , from [Angelopoulos and Kamali, 2023] we know that X must complete a contract C of length $\tau/2$ at time τ , as $\tau \rightarrow \infty$. Consider now the actual interruption $T = \tau - \epsilon$, i.e., the interruption occurs right before C terminates. Then

$$\ell(X, T) = \frac{\tau}{4} = \frac{T + \epsilon}{4}. \quad \square$$

Omitted details of the proof of Theorem 8. In the first part of the proof, we will focus on the effect of ψ^+ . By construction of ψ^+ , we have:

$$\text{EMD}(\mu, \mu^+) = \int_{x \in \mathbb{R}^+} \int_{y > x} (y - x) \psi(x, y) dy dx.$$

For any x , we have $\mathbb{P}_{z \sim \mu^+}(x > z) \geq \mathbb{P}_{z \sim \mu}(x > z)$, therefore we obtain Equation (6) from the main paper:

$$\mathbb{E}_{z \sim \mu^+}[\ell(X(\lambda), z)] \geq \mathbb{E}_{z \sim \mu}[\ell(X(\lambda), z)]. \quad (6)$$

Moreover,

$$\mathbb{E}_{z \sim \mu}[z] = \int_{\mathbb{R}^+} x \cdot \mu(x) dx = \int_{\mathbb{R}^+} \int_{\mathbb{R}^+} x \cdot \psi^+(x, y) dy dx,$$

hence, we obtain Equation (7) from the main paper:

$$\begin{aligned} \mathbb{E}_{z \sim \mu^+}[z] &= \int_{\mathbb{R}^+} \int_{\mathbb{R}^+} y \cdot \psi^+(x, y) dx dy \\ &= \mathbb{E}_{z \sim \mu}[z] + \text{EMD}(\mu, \mu^+) \\ &\leq \mathbb{E}_{z \sim \mu}[z] + \eta. \end{aligned} \quad (7)$$

In the next part of the proof, we focus on the effect of ψ^- . By similar arguments, we can see that, as $\mathbb{P}_{z \sim \mu'}(x > z) \leq \mathbb{P}_{z \sim \mu}(x > z)$, we derive Equation (8) from the main paper:

$$\mathbb{E}_{z \sim \mu'}[z] \leq \mathbb{E}_{z \sim \mu}[z]. \quad (8)$$

We now give a lower bound on $\mathbb{E}_{z \sim \mu'}[\ell(X(\lambda), z)]$. We have

$$\eta \geq \text{EMD}(\mu^+, \mu') = \int_{x \in \mathbb{R}^+} \int_{y < x} (x - y) \psi(x, y) dy dx.$$

Let k such that $2^{k-\lambda} \in [D, 2D)$. As $\mu^+(x) = 0$ for any $x < D$, the last completed contract of $X(\lambda)$ is of length D if the interruption is in the interval $[2^{k-\lambda}, 2D]$ and of length $D/2$ if the interruption is in the interval $(0, 2^{k-\lambda})$. We now give a lower bound on $\mathbb{E}_{z \sim \mu'}[\ell(X(\lambda), z)]$ based on $\mathbb{E}_{z \sim \mu}[\ell(X(\lambda), z)]$. For that, we upper bound, by D , the loss in the potential profit due to probability mass crossing the end of a contract of $X(\lambda)$ because of ψ^- . Indeed, any contract completed by $X(\lambda)$ before time $2D$ has length at most D .

$$\begin{aligned} \mathbb{E}_{z \sim \mu'}[\ell(X(\lambda), z)] &\geq \mathbb{E}_{z \sim \mu}[\ell(X(\lambda), z)] \\ &\quad - D \cdot \int_{y=0}^{2^{k-1-\lambda}} \int_{x=D}^{2^{k-\lambda}} \psi^-(x, y) dx dy \\ &\quad - D \cdot \int_{y=0}^{2^{k-\lambda}} \int_{x=2^{k-\lambda}}^{2D} \psi^-(x, y) dx dy. \end{aligned} \quad (11)$$

Fixing the value of $\text{EMD}(\mu^+, \mu')$, the RHS of (11) is minimized if $\psi^-(x, 2^{k-1-\lambda}) = \mu(x)$ for $x \in [D, D + \delta_1]$ and $\psi^-(x, 2^{k-\lambda}) = \mu(x)$ for $x \in [2^{k-\lambda}, 2^{k-\lambda} + \delta_2]$, for some δ_1, δ_2 , and $\psi^-(x, y) = 0$ for any other $x \neq y$. We now

upper bound δ_1 (resp. δ_2), using $\eta \geq \text{EMD}(\mu^+, \mu')$. Since $\mu(x) \geq \mu(2D) = 1/(2D)$, we have:

$$\begin{aligned} \eta &\geq \int_D^{D+\delta_1} |x - 2^{k-1-\lambda}| \cdot \psi^-(x, 2^{k-1-\lambda}) dx \\ &\geq \int_{D+\delta_1/2}^{D+\delta_1} \frac{\delta_1}{2} \cdot \mu(2D) dx \\ &\geq \frac{\delta_1}{2} \cdot \frac{\delta_1}{2} \cdot \frac{1}{2D} \\ \delta_1 &\leq \sqrt{8\eta D}. \end{aligned} \quad (12)$$

We have the same result for δ_2 , so both δ_1 and δ_2 are smaller than $\sqrt{8\eta \cdot D}$.

We now upper bound the decrease in $\mathbb{E}_{z \sim \mu'}[\ell(X(\lambda), z)]$ that follows from Equation (12). We can bound the integral:

$$\begin{aligned} \int_{2^{k-1-\lambda}}^{2^{k-1-\lambda}+\delta_1} \mu(x) &\leq \delta_1 \cdot \mu(D) \\ &\leq \frac{2}{D} \cdot \sqrt{8\eta \cdot D} \\ &\leq 4\sqrt{\frac{2\eta}{D}}, \end{aligned}$$

and equivalently for δ_2 . Therefore, using Equation (11), we obtain Equation (9) from the main paper:

$$\begin{aligned} \mathbb{E}_{z \sim \mu'}[\ell(X(\lambda), z)] &\geq \mathbb{E}_{z \sim \mu^+}[\ell(X(\lambda), z)] - 2D \cdot 4\sqrt{\frac{2\eta}{D}} \\ &\geq \mathbb{E}_{z \sim \mu^+}[\ell(X(\lambda), z)] - 8\sqrt{2\eta D}. \end{aligned} \quad (9)$$

To conclude the proof, note that $\mathbb{E}_{z \sim \mu}[z] = 2D \ln 2$ and $\mathbb{E}_{z \sim \mu}[\ell(X(\lambda), z)] = D/2$, from the proof of Theorem 6. Combing Equations (6), (7), (8) and (9), we obtain for $\eta < D/512$ (to ensure there is no division by zero):

$$\begin{aligned} \frac{\mathbb{E}_{z \sim \mu'}[z]}{\mathbb{E}_{z \sim \mu'}[\ell(X(\lambda), z)]} &\leq \frac{\mathbb{E}_{z \sim \mu^+}[z]}{\mathbb{E}_{z \sim \mu^+}[\ell(X(\lambda), z)] - 8\sqrt{2\eta D}} \\ &\leq \frac{\mathbb{E}_{z \sim \mu}[z] + \eta}{\mathbb{E}_{z \sim \mu}[\ell(X(\lambda), z)] - 8\sqrt{2\eta D}} \\ &\leq \frac{2D \ln 2 + \eta}{\frac{D}{2} - 8\sqrt{2\eta D}} \\ &\leq \frac{4 \ln 2 + 2\eta/D}{1 - 16\sqrt{2\eta/D}}. \end{aligned}$$

Finally, for small values of η/D , we have:

$$\frac{\mathbb{E}_{z \sim \mu'}[z]}{\mathbb{E}_{z \sim \mu'}[\ell(X(\lambda), z)]} \leq 4 \ln 2 + O(\sqrt{\eta/D}),$$

which concludes the proof. \square

B Additional experimental results

We report additional experimental results concerning the distributional advice model.

B.1 Time horizon

We repeated the experiments of Section 5.1 for larger time horizons, namely $[1, 20000]$, shown in Figure 6 and Figure 7. The plots have the same overall shape as those reported in the main paper, which is expected, considering that the schedule follows a scaled doubling strategy.

B.2 Fixed variance

We considered again the setting of Section 5.1, with the difference that the variance of the distributional advice is now fixed, instead of being a function of time (or the mean m). This captures the situation in which the window of time in which we expect the interruption to occur is prespecified. Specifically, Figure 8 depicts the consistency of SEL_n given a normal distributional advice of mean m , and standard deviation $\sigma = 10$.

We observe that as the time increases, the transitions of the plot become sharper, relatively to the setting in which the variance is fixed. This is because, as the time increases, the variance becomes less significant in relation to the mean of the distributional advice. This implies, in turn, that the advice has decreasing spread as function of time, and becomes more and more close to a single-valued prediction for large t . This also explains why for the same value of n , the consistency of SEL_n is worse in this setting than in the setting of Figure 3, in the main paper. The values reported for very small values of m are outlier, and are due to the high variance (relative to the small values of m).

B.3 Time-dependent variance

We did experiments with a distributional advice whose mean is variable, and function of time, and whose variance is larger or smaller than the values of the main paper. Figure 9 and Figure 10 depict the experimental consistency with normal distributional advice for a setting as in the main paper, and standard deviation $\sigma = 0.01m$ and $\sigma = 0.2m$, respectively. For small σ , we observe sharper transitions, and a larger gap between the peaks and valleys, which can be explained along the same lines as in the discussion of Section B.2. For large σ , we observe that the consistency increases to values close to 2.7, and that the plots become relatively “flat”, with much smoother transitions. This is because, with large variance, the setting becomes highly random, and the prediction less crucial or useful, which also explains the much smaller gap between the various values of n .

Similar observations and conclusions can be drawn for uniform advice, as depicted in Figure 11.

B.4 Error analysis

We conducted experiments to test the resilience of the algorithm SEL_n to errors in the distributional advice. In particular, we considered the setting in which the advice is a normal distribution with mean m and standard deviation fixed to $\sigma = 25$, whereas the interruption time is drawn according to a normal distribution with *actual*, but unknown mean m' (in general, $m' \neq m$) and standard deviation equal to 25. We choose $n = 16$.




Figure 6: Plot of the consistency of SEL_n with time horizon $[1, 20000]$, for normal distributional advice as in the main paper.




Figure 7: Plot of the consistency of SEL_n with time horizon $[1, 20000]$, for uniform distributional advice as in the main paper.




Figure 8: Plot of the consistency of SEL_n for normal distributional advice with mean m and fixed standard deviation $\sigma = 10$.




Figure 9: Plot of the consistency of SEL_n for normal distributional advice with mean m and standard deviation $\sigma = 0.01m$.

In Figure 12, the blue curve depicts the performance of our schedule as function of the actual mean m' (and thus, as function of the error, defined as $\eta = |m' - m|$). In contrast, the red curve depicts the consistency of the schedule with error-free advice. We observe that as long as the error is relatively small, the schedule remains fairly robust to prediction errors, and shows smooth degradation. We also observe an asymmetry on the performance, depending on whether the error is “positive”, i.e., $m' > m$ or “negative”, i.e., $m' < m$). This is expected, since a positive error does not have as critical an effect as negative error: positive error implies that a large contract in the schedule will still likely terminate, whereas negative error implies that the schedule is running a large contract which likely will *not* terminate by the interruption.

Figure 13 depicts the same plots for much larger values of error. We note that the performance plot exhibits, informally, a certain periodicity, which can be explained by the nature of the problem. Namely, even if the error becomes large, it is possible that the advice may have a beneficial effect to the schedule, if the latter happens to complete a large contract

close to the actual, expected interruption time. This explains why the blue and red curves meet not only for $m' = m$, but also for other values of m' . This finding is also an unavoidable situation: *any* schedule will perform very well for some outlier values of large error, which implies that it is impossible to find a schedule that strictly dominates all other schedules across any error domain.

Figure 14 depicts the same plot for normal distributional advice with $m = 700$. We observe similar performance and findings as discussed above.




Figure 10: Plot of the consistency of SEL_n for normal distributional advice with mean m and standard deviation $\sigma = 0.2m$.




Figure 11: Plot of the consistency of SEL_n for uniform distributional advice $U[0.9t, 1.1t]$, as function of time t .




Figure 12: Plot of the consistency of SEL_{16} with normal distributional advice of mean $m = 500$, as function of the actual mean m' .




Figure 13: Plot of the consistency of SEL_{16} with normal distributional advice of mean $m = 500$, as function of the actual mean m' (larger range of error).




Figure 14: Plot of the consistency of SEL_{16} with normal distributional advice of mean $m = 700$, as function of the actual mean m' (larger range of error).