



**HAL**  
open science

## Competitive Sequencing with Query Predictions

Spyros Angelopoulos, Diogo Arsénio, Shahin Kamali

► **To cite this version:**

Spyros Angelopoulos, Diogo Arsénio, Shahin Kamali. Competitive Sequencing with Query Predictions. 2024. hal-04713339

**HAL Id: hal-04713339**

**<https://hal.science/hal-04713339v1>**

Preprint submitted on 29 Sep 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.

# Competitive Sequencing with Query Predictions<sup>\*</sup>

Spyros Angelopoulos<sup>a,\*</sup>, Diogo Arsénio<sup>b</sup>, Shahin Kamali<sup>c</sup>

<sup>a</sup>*Sorbonne University, CNRS, LIP6, Paris, France*

<sup>b</sup>*New York University, Abu Dhabi, UAE*

<sup>c</sup>*Department of Electrical Engineering and Computer Science, York University, Toronto, Canada*

---

## Abstract

Several well-studied online resource allocation problems can be abstracted in terms of an infinite, increasing sequence, where each element is associated with a corresponding allocation value. In Theoretical Computer Science, one such abstraction is known as *online bidding*, in which an algorithm must submit “bids” until an unknown threshold is reached. Another abstraction that has been studied extensively in Artificial Intelligence is known as *contract scheduling*: in this formulation, an algorithm is repeatedly executed with increasing processing times so as to obtain a system with interruptible capabilities.

We study such problems under the *query prediction* model, in which the designer elicits a prediction on the instance via responses to  $k$  binary queries so as to improve the algorithm’s performance. The queries are answered by an imperfect oracle, and the objective is to obtain efficient algorithms

---

<sup>\*</sup>Partially supported by the project PREDICTIONS, grant ANR-23-CE48-0010 from the French Research Agency (ANR), and Natural Sciences and Engineering Research Council of Canada (NSERC), grant DGEER-2018-00059.

<sup>\*</sup>Corresponding author

*Email addresses:* `spyros.angelopoulos@lip6.fr` (Spyros Angelopoulos), `diogo.arsenio@nyu.edu` (Diogo Arsénio), `kamalis@yorku.ca` (Shahin Kamali)

that are also robust to query errors. We first focus on *consistency-robustness* tradeoffs, in which the query responses are either error-free or are generated by a (malicious) adversary: here, we prove a tight information-theoretic lower bound that establishes *Pareto-optimality* with respect to the consistency and the robustness. Next, we consider the more general setting in which some of the query responses can be erroneous: here, we present and analyze an efficient and robust algorithm based on *adaptive* queries. Specifically, we show that small increments in the number of queries lead to substantial improvement in robustness to query errors, in that the performance of our solution approaches the ideal performance of the Pareto-optimal schedule very quickly as  $k$  grows, even if as many as  $k/4 - o(k)$  responses are adversarially erroneous. Our techniques have applications outside the query model: Namely, we show how to obtain an optimal schedule for a generalization of the *fault tolerant* contract scheduling problem in a multi-processor system, which generalizes the setting of [Kupavskii and Welzl, *Distr. Comp.* 2019].

*Keywords:* Competitive analysis, algorithms with predictions, consistency, robustness, query models.

---

## 1. Introduction

Resilience to interruptions is a central requirement in the design of real-time and intelligence systems. For instance, applications such as medical diagnostic systems, robot motion planning, and financial planning and trading require that the system be capable of outputting a reasonably efficient solution at all times [1]. This raises the important issue of designing an interruptible system given, as a building component, an algorithm that does

8 not necessarily have interruptible capabilities.

9 The seminal work of Russell and Zilberstein [2] was the first to provide a  
10 methodology for the design of interruptible systems using *contract* algorithms  
11 as components. The latter are algorithms which, unlike the interruptible  
12 ones, require the amount of computational time to be known in advance.  
13 Specifically, if a contract algorithm is allowed an execution time that is at  
14 least as large as this “contract” time, then its output is guaranteed to be  
15 correct; however, if the algorithm is interrupted at any point prior to the  
16 contract time, then its output may be totally meaningless. A typical ex-  
17 ample is algorithms based on dynamic programming (DP); if the algorithm  
18 fails to fill the entire DP table, the output may be entirely useless. Despite  
19 this lack of flexibility, contract algorithms are often easier to implement and  
20 maintain and use relatively simpler data structures [3], which makes them of-  
21 ten excellent components for the design of more complex, real-time systems.  
22 Note that interruptible and contract algorithms are members of the broader  
23 class of *anytime* algorithms, i.e., algorithms whose performance improves as  
24 a function of the available computational time.

25 The approach of [2] is based on iterative deepening and consists of re-  
26 peatedly executing the contract algorithm with increasing contract times  
27 (also called *lengths*). To illustrate with an example, consider a simple dou-  
28 bling rule in which the  $i$ -th execution has length  $2^i$ . In this case, even if  
29 an interruption occurs at a worst-case time instance  $T$  (namely, right before  
30 an execution is about to terminate), an execution of length at least  $T/4$  has  
31 been completed. The factor 4 measures the performance of this doubling rule  
32 and describes the tradeoff between resilience to interruptions and processor

33 speed. That is, a system  $A$  based on doubling contract lengths, and with  
 34 processor speed equal to 4, is at least as efficient as any system  $B$  of unit  
 35 speed, even if  $B$  knows beforehand the interruption time (and could thus use  
 36 only a single execution with contract time  $T$ ).

37 More generally, given a contract algorithm  $A$ , a contract *schedule* can  
 38 be described as an increasing sequence  $(x_i)_{i \geq 0}$ , in which  $x_i$  is the length of  
 39 the  $i$ -th execution of  $A$ . The work [2] introduced a worst-case, theoretical  
 40 measure, akin to the competitive ratio, for evaluating the performance of  $X$ ,  
 41 which relates an interruption time  $T$  to the largest contract length completed  
 42 by time  $T$ . This measure is called the *acceleration ratio* of  $X$ . Formally,

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

43 where  $\ell(X, T)$  denotes the length of the largest contract completed in  $X$  by  
 44 time  $T$ .

45 Contract scheduling is a classic resource-allocation problem that has been  
 46 studied in many settings. For the standard version described above, it is  
 47 straightforward to show that the doubling rule yields a schedule of optimal  
 48 acceleration ratio equal to 4, however the problem becomes substantially  
 49 more involved under more complex settings. We review some related results.  
 50 Schedules for multi-processor systems were first obtained in [4]. The work [5]  
 51 studied the setting in which there are several problem instances that must  
 52 be solved concurrently on a single processor, and [3] as well as [6] provided  
 53 optimal acceleration ratios for the multi-instance/multi-processor generaliza-  
 54 tion. A study of the setting in which the interruption is not a hard deadline  
 55 (and the system is allowed some additional time to wrap up its execution was  
 56 given in [7], whereas [8] introduced performance measures alternative to the

57 acceleration ratio. [9] considered a setting in which the schedule is deemed  
 58 complete once a length-related guarantee is reached on its last completed  
 59 contract. Connections between contract scheduling and searching for a hid-  
 60 den target under the competitive ratio were shown in [3, 10]. We emphasize  
 61 that all of the above works establish theoretical upper and lower bounds on  
 62 the acceleration ratio.

63 Contract scheduling can be considered as an application of a general *se-*  
 64 *quencing* formulation, in which we seek an increasing sequence that optimizes  
 65 a given performance measure. Another sequencing problem that has been  
 66 well studied within TCS is *online bidding* [11, 12, 13, 14, 15, 16]. Here, the  
 67 objective is to find an increasing sequence  $X = (x_i)_{i \geq 0}$  of positive numbers  
 68 of minimum *competitive ratio*, defined formally as

$$\sup_{u \geq 0} \frac{\sum_{j=1}^i x_j}{u} : x_{i-1} < u \leq x_i, \quad (2)$$

69 where  $u$  represents an unknown *target* value. Online bidding has applications  
 70 in problems such as clustering [12], searching on the infinite line [17] and  
 71 latency minimization [18], see e.g., the survey [11]. The doubling strategy  
 72  $X = (2^i)_{i \geq 0}$  achieves once again an optimal competitive ratio equal to 4.

### 73 1.1. Competitive sequencing with queries

74 Recently, contract scheduling was studied under a model in which the  
 75 scheduler leverages predictions on the interruption time, in the form of re-  
 76 sponses to  $k$  *binary queries*, for some given  $k \in \mathbb{N}^+$  [19]. For example, a query  
 77 may be of the form “will the interruption occur before time  $t = 100$ ?”, or may  
 78 be more complex, e.g., “will the interruption occur in  $\cup_{i \text{ odd}} [2^i, 2^{i+1}]$ ?”. For-  
 79 mally, a query maps a statement concerning the interruption time to  $\{0, 1\}$ .

80 As explained in [19], queries can help improve resource planning in many  
 81 settings. For instance, in a medical-diagnostic system, the end user may  
 82 know that the system is likely to be consulted at some critical intervals, e.g.,  
 83 around some anticipated surgery slots, or that it may be more likely that a  
 84 diagnosis would be needed on a weekday rather than on a weekend.

85 Figure 1 illustrates an example for the simple case of a single query. Note  
 86 that in the context of sequencing problems such as contract scheduling and  
 87 online bidding, the query responses are powerful enough to encode *interval*  
 88 information about the interruption or the target. Such information cannot  
 89 be captured by simpler prediction oracles that provide a single value that  
 90 corresponds to the anticipated interruption or target. The queries are an-  
 91 swered by an imperfect oracle; namely,  $\eta \leq k$  queries may receive erroneous  
 92 responses. We refer to  $\eta$  as the *query error*, and we note that  $\eta$  is not known  
 93 to the system designer ahead of time.




Figure 1: An illustration of the query-based setting for  $k = 1$ . Here, a single query may answer whether the interruption is in the yellow (lighter) or in the blue (darker) partition. The decision-maker may use the query response to choose one of the schedules  $S_1$  and  $S_2$ .

94 The query prediction model combines two essential aspects of *learning-*  
 95 *augmented* optimization: The first aspect is the use of *predictions* in op-  
 96 timization under uncertainty, in which the (typically online) algorithm is  
 97 augmented with an inherently erroneous prediction oracle, and the objec-

98 tive is to improve the algorithm’s performance via a robust leveraging of the  
99 prediction oracle. Such considerations have become very prevalent in recent  
100 studies of algorithmic performance, starting with the influential works [20]  
101 and [21], see e.g., the survey [22] and the online repository [23] that lists  
102 several related works. The second aspect pertains to *query-based* optimiza-  
103 tion, in which the algorithm recovers the solution to a problem by asking  
104 queries. An example is clustering with noisy queries [24, 25], where a query  
105 asks whether two elements belong in the same cluster, which is useful in  
106 crowdsourcing applications. Query-based oracles can thus help model *parsi-*  
107 *monious* predictions, and similar models were recently studied in the context  
108 of secretary problems [26] and online paging [27].

109 As in the analysis of learning-augmented algorithms, e.g. [28, 29], we first  
110 evaluate the performance of the algorithms in two extreme situations with  
111 respect to the prediction error. In the one extreme, the oracle is perfect,  
112 thus  $\eta = 0$ : in this case, we refer to the acceleration ratio of a schedule as its  
113 *consistency*. In the other extreme, the oracle is malicious (i.e., all responses  
114 are adversarial): in this case, we refer to the acceleration ratio of a schedule  
115 as its *robustness*. Here, the goal is to find *Pareto-optimal* schedules that  
116 describe the optimal consistency/robustness tradeoffs, similarly for online  
117 bidding under the competitive ratio. Pareto-efficient algorithms with respect  
118 to consistency-robustness tradeoffs have become prominent in the context of  
119 online optimization problems with untrusted predictions, see e.g., [14, 28, 29,  
120 30, 31, 32, 33, 34].

121 Beyond these two extremes, we are also interested in efficient algorithms  
122 that are robust to query errors. To this end, [19] introduced a *noisy query*



123 *model*, specified as follows. The decision maker defines a parameter  $\tau \leq k$ ,  
 124 which is interpreted as an anticipated *upper bound* on the query error or, al-  
 125 ternatively, as the algorithm’s desired *tolerance* to the error. This parameter  
 126 captures the requirement that the algorithm must perform well if this upper  
 127 bound is met, i.e., if  $\eta \leq \tau$ , but must also have robustness at most  $r$ , for some  
 128 specified  $r \geq 4$ , if it so happens that  $\eta > \tau$ . Such a parameter is common  
 129 in the analysis of *games with a lying responder*, e.g., [35], in which an upper  
 130 bound on the erroneous responses is assumed to be known. It is also related  
 131 to the concept of *weak predictions*, (see, e.g., online knapsack with frequency  
 132 predictions [36], in which the prediction is an upper bound on the number  
 133 of items of a given value in the input). We emphasize that unlike the noisy  
 134 query model of [24], we do not make any probabilistic assumptions on the  
 135 query responses.

## 136 1.2. Contribution

137 **Results.** Our first result (Theorem 13) proves the Pareto-optimal trade-off  
 138 between the consistency  $c$  and the robustness  $r$  for contract scheduling with  
 139  $k$  queries, for any value of  $r$ . This answers the main problem left open in [19],  
 140 which showed a tight tradeoff only for the special case of  $r = 4$ . To prove our  
 141 result, we give a tight, information-theoretic lower bound on the consistency  
 142 of  $r$ -robust schedules, which matches the known upper bound (Theorem 11).  
 143 This result also establishes properties useful in the other settings we study.  
 144 We also show, using a reduction from contract scheduling to online bidding,  
 145 that the same optimal guarantees carry over to the latter problem as well  
 146 (Theorem 15). This result answers an open question from [14] that gave a  
 147 non-tight lower bound assuming  $r = 4$ .

148 Our second result concerns the general query model. Here, given a ro-  
 149 bustness requirement  $r$ , we obtain a schedule that is optimal in the ideal  
 150 case that  $\eta = 0$ , remains  $r$ -robust for adversarial error (i.e., if  $\eta$  is as high  
 151 as  $k$ ), and has acceleration ratio that provably degrades gently as a function  
 152 of the anticipated bound on the error  $\tau$ , as long as  $\eta \leq \tau \leq k/4 - o(k)$ <sup>1</sup>  
 153 (Theorem 20). More importantly, we show that as long as the error is not  
 154 prohibitively large, namely for any  $\eta \leq k/4 - o(k)$ , the acceleration ratio  
 155 converges very quickly to the error-free consistency as  $k$  increases. In other  
 156 words, we prove that small increments in the number of queries lead to signif-  
 157 icant gains in robustness to errors, even for substantially and unpredictably  
 158 erroneous queries. This is in contrast to [19], where the acceleration ratio  
 159 converges to a much larger and sub-optimal value (Observation 21). Our  
 160 analysis also implies that the scheduler is not constrained by the choice of a  
 161 particular tolerance parameter  $\tau$ , even if  $k$  is a rather small constant. The  
 162 same guarantees can be extended to online bidding.

163 To prove the above results, we develop techniques that are also applicable  
 164 to multi-criteria contract scheduling unrelated to the query setting. As an  
 165 example, in our third main result, we introduce and study a *robust* general-  
 166 ization of *parallel, fault-tolerant* contract scheduling. In the original problem  
 167 studied by Kupavskii and Welzl [37], the objective is to optimize the acceler-  
 168 ation ratio of a contract schedule in a parallel  $p$ -processor system, assuming  
 169 that at most  $f$  processors may be faulty, for some given  $f < p$ ; however, the  
 170 schedule may be arbitrarily bad if the number of faulty processors exceeds  $f$ .

---

<sup>1</sup>Here, the term  $o(k)$  describes any fixed, and slowly increasing function  $f(k) \in o(k)$ ,  
 e.g.,  $f(k) = \log^*(k)$ .

171 In our setting, instead, we additionally require that the system remains effi-  
172 cient if all but a single processor are faulty. We show how to obtain a schedule  
173 with the best-possible tradeoff between the two performance objectives (The-  
174 orems 22 and 24). Last, in Section 6 we provide an experimental evaluation  
175 of our algorithms that demonstrates the performance improvements that are  
176 attained in practice.

177 **Techniques.** In regards to Pareto-optimality, we give a lower bound on the  
178 consistency of any  $r$ -robust schedule with  $k$  queries by treating it as a “vir-  
179 tual” multi-processor scheduling problem in  $2^k$  identical parallel processors,  
180 where each processor executes an  $r$ -robust schedule. The robustness require-  
181 ment adds significant complexity to our setting, which is more involved than  
182 the standard multi-processor contract scheduling studied by López-Ortiz et  
183 al. [6]. Specifically, while the algorithm and the analysis in [6] allow for  
184 schedules that are non-robust, in our setting, each processor is required to  
185 implement an  $r$ -robust schedule. This necessitates a more careful analysis by  
186 looking deeper into the linear recurrence relations (equalities and inequali-  
187 ties) that formulate the concept of robustness, as we will show in Section 3.  
188 In particular, in order to establish Pareto-optimality, we relate inequalities  
189 expressed via linear recurrences to the solution of linear recurrence relations  
190 (e.g., in the statement and proof of Proposition 6) and rely on upper-limit  
191 calculus (e.g., in the proof of Lemma 9).

192 To achieve robustness to query errors, we first define a space of  $2^k$  sched-  
193 ules that include the Pareto-optimal one and which has a “nice” structure:  
194 namely, there exists an ordering such that if the  $j$ -th schedule has the best  
195 performance (for a given interruption), then we can bound the acceleration

196 ratio of the  $(j - x)$ -th schedule in the ordering as a function of  $x$ . We then use  
 197 an error-tolerant binary search algorithm inspired by Disser and Kratsch [38]  
 198 based on the query responses, so as to find a schedule close to the best one,  
 199 even in the presence of errors, and without any knowledge of the ordering.

200 We emphasize that our approach is based on *adaptive* queries, in that the  
 201  $i$ -th query is a function of the responses to queries  $1 \dots i - 1$ . Adaptive queries  
 202 allow searching an exponential space of candidate schedules, unlike [19] which  
 203 relies on static queries that can only help search a linear space of schedules (as  
 204 a function of the number of queries  $k$ ). We thus demonstrate that adaptivity  
 205 is important for optimality.

## 206 2. Preliminaries

207 We introduce notation for the contract scheduling problem; for online  
 208 bidding, our results carry over via a reduction shown in Theorem 15. In a  
 209 single-processor system, a schedule  $X$  is defined as an increasing sequence  
 210 of the form  $X = (x_i)_{i \geq 0}$ . We make the standing assumption that the inter-  
 211 ruption occurs after at least the first contract has completed its execution;  
 212 otherwise, no schedule has a bounded acceleration ratio. Without queries, the  
 213 acceleration ratio of  $X$  is given by (1). It is well-known that the worst-case  
 214 interruptions that maximize  $\text{acc}(X)$  are infinitesimally prior to the comple-  
 215 tion times of contracts, namely  $x_i - \epsilon$ , for  $\epsilon \rightarrow 0$ . Hence the equivalent  
 216 expression (where  $x_{-1}$  is defined to be equal to 1)

$$\text{acc}(X) = \sup_i \frac{\sum_{j=0}^i x_j}{x_{i-1}}. \quad (3)$$

217 For a schedule  $X$  with  $k$  queries, [19] defines the *consistency* of  $X$  as

218 its acceleration ratio assuming no response errors, and its *robustness* as its  
 219 acceleration ratio assuming worst-case, adversarial responses. A schedule is  
 220 *r-robust* if its robustness is at most  $r$  (similarly for consistency), and it is  
 221 *Pareto-optimal* if its consistency and robustness are in a best-possible rela-  
 222 tion. Note that the robustness of a schedule is equal to its acceleration ratio  
 223 in the standard setting (since both definitions involve worst-case, adversarial  
 224 settings). Hence, it is always the case that  $r \geq 4$ , and we will thus use these  
 225 two terms interchangeably.

226 A schedule  $G_b$  is called *geometric with base  $b > 1$*  if it is of the form  
 227  $G_b = (b^i)_{i \geq 0}$ . Geometric schedules are significant since they are often efficient  
 228 for several variants of contract scheduling. It is known that the acceleration  
 229 ratio of  $G_b$  is equal to  $b^2/(b-1)$  [5], which is minimized for  $b = 2$ , thus  
 230  $G_2 = (2^i)_{i \geq 0}$  has optimal acceleration ratio equal to 4. More generally, one  
 231 can easily identify the geometric schedules that are *r-robust*, for any given  $r$ .

232 **Definition 1.** For any given  $r$ , define  $\zeta_{1,r}$  and  $\zeta_{2,r}$  as the smallest and largest  
 233 real roots, respectively, of the function  $f(x) = \frac{x^2}{x-1} - r$ . It also holds that  
 234  $1 < \zeta_{1,r} \leq \zeta_{2,r}$ .

235 This definition implies the following useful property:

236 **Property 2.** The schedule  $G_b$  with  $b > 1$  has acceleration ratio at most  
 237  $b^2/(b-1)$ . In particular, for any  $r \geq 4$ ,  $G_b$  has acceleration ratio at most  $r$   
 238 if and only if  $b \in [\zeta_{1,r}, \zeta_{2,r}]$ .

239 The roots of Definition 1 are easily computable, i.e.,  $\zeta_{1,r} = (r - \sqrt{r^2 - 4r})/2$   
 240 and  $\zeta_{2,r} = (r + \sqrt{r^2 - 4r})/2$ , and are such that  $\zeta_{1,r} \in (1, 2]$ , and  $\zeta_{2,r} \geq 2$ , for  
 241 any  $r \geq 4$ .

242 **Example 3.** Suppose that  $r = 6$ , then  $\zeta_{1,6} = 3 - \sqrt{3}$ , and  $\zeta_{2,6} = 3 + \sqrt{3}$ . Any  
 243 geometric schedule  $G_b$ , with  $b \in [3 - \sqrt{3}, 3 + \sqrt{3}]$  is 6-robust.

244 The above definitions assume a single processor. In our work, we show  
 245 and exploit connections between single-processor schedules with  $k$  queries and  
 246 multi-processor schedules without queries on  $2^k$  parallel processors. Hence,  
 247 we present some definitions and notation concerning the setting of  $p > 1$   
 248 parallel processors, labelled  $\{0, \dots, p - 1\}$ . In the  $p$ -processor setting, each  
 249 processor  $j$  defines its own *strategy* of the form  $X_j = (x_{i,j})_{i \geq 0}$ . We call the  
 250 set  $X = \{X_j\}_{j=0}^{p-1}$  a  $p$ -processor schedule, or equivalently, we say that  $X$  is  
 251 defined by the set  $\{X_j\}_{j=0}^{p-1}$ . The acceleration ratio of a  $p$ -processor schedule is  
 252 defined as in (1), with the difference that  $\ell(X, T)$  denotes the largest contract  
 253 length completed by time  $T$  among all  $p$  processors [3].

254 Let  $Y = (y_i)_{i=0}^{\infty}$  denote a positive sequence, and define  $\alpha_Y = \limsup_{n \rightarrow \infty} y_n^{1/n}$ .  
 255 E.g., if  $Y = (2^i)_{i=0}^{\infty}$ , then  $\alpha_Y = 2$ . This notion appears in the statement of  
 256 a theorem by Gal [39], which is the basis of the analysis in the multipro-  
 257 cessor setting of [6]. This theorem, informally, gives a lower bound on the  
 258 supremum of a set of functionals by the supremum of these functionals over  
 259 geometrically increasing sequences. Given a set (or sequence)  $Y$  of positive  
 260 reals, we denote by  $\bar{Y}$  the sequence of all elements in  $Y$  in non-decreasing  
 261 order. We refer to Appendix Appendix A for the formal statement of Gal's  
 262 theorem, which will not be of direct use in our work, though the notion of  $\alpha_Y$   
 263 will figure prominently. Table 1 summarizes some important notation used  
 264 throughout this work.

Notation	Description
$T$	interruption time (unknown to the scheduler)
$\ell(X, T)$	largest contract completed in schedule $X$ by time $T$
$\text{acc}(X)$	acceleration ratio of schedule $X$
$k$	number of queries
$G_b$	geometric strategy with base $b > 1$ , i.e., of the form $(b^i)_{i=0}^\infty$
$r$	robustness ( $r \geq 4$ )
$\zeta_{1,r}, \bar{\zeta}_{1,r}$	the smallest and largest positive roots of the function $f(x) = \frac{x^2}{x-1} - r$ .
$\alpha_Y$	$\limsup_{n \rightarrow \infty} y_n^{1/n}$ of the sequence $Y = (y_i)_{i=0}^\infty$
$\mathcal{X}_{b,l}$	the set of strategies $\{X_0, \dots, X_{l-1}\}$ , where $X_i = (b^{i+jl})_{j=0}^\infty$ ( $i \in [0, 2^k - 1]$ ) (see Definition 12.)
$\eta$	(unknown) number of erroneous query responses ( $\eta \leq k$ )
$\tau$	anticipated upper bound on $\eta$ , or tolerance
$U$	defined as $2^{\lfloor k/2 \rfloor + 2\tau}$
$c(k, r)$	consistency of the Pareto-optimal schedule of Theorem 11
$R(k, r, \tau)$	acceleration ratio of the schedule RQS of Theorem 18

Table 1: Summary of notation and definitions.

### 265 3. Pareto-optimality of contract scheduling and online bidding

266 In this section, we give the optimal consistency-robustness tradeoff for  
267 both sequencing problems, with access to  $k$  queries. Our main result is a  
268 tight information-theoretic lower bound that applies to any schedule with  
269 information encoded as a  $k$ -bit string, which we call the *response string*.

270 **Overview of the proof.** We first give an overview and the intuition of  
 271 the proof. The starting observation is that any  $r$ -robust schedule  $X$  with  $k$   
 272 queries must be selected from a set  $\mathcal{X}$  of at most  $2^k$   $r$ -robust schedules. This  
 273 implies that the consistency of  $X$  is the acceleration ratio of the  $2^k$ -processor  
 274 schedule defined by  $\mathcal{X}$ . We can use a lower bound on this acceleration ratio as  
 275 a function of the parameter  $\alpha_{\bar{\mathcal{X}}}$  (Lemma 4), and recall that  $\bar{\mathcal{X}}$  is the sequence  
 276 of all contract lengths of schedules in  $\mathcal{X}$ , in non-decreasing order.

277 There is, however, a substantial complication, in that it is imperative to  
 278 show that  $\alpha_{\bar{\mathcal{X}}}$  is within a certain range in order to establish the tightness of  
 279 the lower bound; this is to capture the requirement that each strategy in  $\mathcal{X}$   
 280 must be  $r$ -robust. To this end, we show that  $\alpha_{\bar{\mathcal{X}}} \in [\zeta_{1,r}^{1/2^k}, \zeta_{2,r}^{1/2^k}]$  (Corollary 10).  
 281 This is accomplished by first showing upper and lower bounds on the contract  
 282 lengths of any  $r$ -robust schedule (Theorem 5), then applying the definitions of  
 283 the sorted sequence of contract lengths  $\bar{\mathcal{X}}$  and the properties of upper limits  
 284 (Lemma 9). Combining the above yields the lower bound (Theorem 11). The  
 285 tightness of the result will follow by directly comparing to the upper bound  
 286 of [19] (Theorem 13).

287 **Analysis.** We now proceed with the technical analysis. The following lemma  
 288 is due to [6] and is a special case of a more general result that incorporates  
 289 fault tolerance and which we will prove later, namely Lemma 23.

290 **Lemma 4.** [6] *Let  $X$  be a  $p$ -processor schedule, as defined by a set  $\mathcal{X}$  of  $p$*   
 291 *single-processor strategies, each of which has a finite acceleration ratio. Then*

$$292 \text{acc}(X) \geq \frac{\alpha_{\bar{\mathcal{X}}}^{p+1}}{\alpha_{\bar{\mathcal{X}}}^p - 1}.$$

293 In the next step, we show upper and lower bounds on the contract lengths  
 294 of any  $r$ -robust schedule. We emphasize that these bounds apply to any



295 schedule of acceleration ratio  $r$  without queries. This will be instrumental in  
 296 bounding the range of  $\alpha_{\bar{\chi}}$ .

297 **Theorem 5.** *Let  $Z = (z_i)_{i \geq 0}$  be an  $r$ -robust single-processor schedule, for a  
 298 fixed, finite  $r \geq 4$ . Then there exist  $c, d$  which are only functions of  $r$ , such  
 299 that  $z_i \leq c \cdot \zeta_{2,r}^i$  and  $z_i \geq d \cdot \zeta_{1,r}^i$ , if  $r > 4$ . Moreover,  $z_i \leq c \cdot i \cdot 2^i$  and  $z_i \geq d \cdot 2^i$ ,  
 300 if  $r = 4$ .*

301 To prove Theorem 5, we first show the following technical result that  
 302 relies on linear recurrence relations. For some intuition, Proposition 6 relates  
 303 a general  $r$ -robust strategy  $X$  to another  $r$ -robust strategy  $Y$ , which satisfies  
 304 all constraints with equality. This will be helpful not only in establishing the  
 305 upper bounds in Theorem 5 but, more importantly, the lower bounds, which  
 306 is the harder part.

307 **Proposition 6.** *Let  $(x_i)_{i=0}^\infty$  and  $(y_i)_{i=0}^\infty$  be sequences of positive numbers such  
 308 that*

$$\sum_{i=0}^n x_i \leq r x_{n-1} \quad \text{and} \quad \sum_{i=0}^n y_i = r y_{n-1}, \quad (4)$$

309 *for all  $n \geq 2$ , where  $r \geq 4$  is a fixed constant. If  $x_0 \geq y_0$  and  $x_1 \leq y_1$ , then  
 310  $x_i \leq y_i$ , for all  $i \geq 2$ .*

311 *Proof.* First, we introduce the sequence of coefficients  $(A_t, B_t)_{t=1}^\infty$  defined  
 312 recursively by

$$\begin{pmatrix} A_{t+1} \\ B_{t+1} \end{pmatrix} = M \begin{pmatrix} A_t \\ B_t \end{pmatrix},$$

313 where

$$M = \begin{pmatrix} r-1 & -1 \\ 1 & 1 \end{pmatrix},$$

314 with initial values  $(A_1, B_1) = (r - 1, 1)$ . We claim that

$$A_t, B_t \geq 0, \quad \text{for every } t \geq 1. \quad (5)$$

315 Let us momentarily assume the validity of (5) and complete the proof of the  
316 proposition. To that end, we show that

$$\begin{aligned} x_n &\leq A_t x_{n-t} - B_t \sum_{i=0}^{n-t-1} x_i \\ \text{and } y_n &= A_t y_{n-t} - B_t \sum_{i=0}^{n-t-1} y_i, \end{aligned} \quad (6)$$

317 for all  $1 \leq t < n$ . This follows from an induction argument on  $t$ . Indeed, for  
318 a given  $n \geq 2$ , the case  $t = 1$  is a mere reformulation of (4). Then, in view  
319 of (5), assuming that (6) holds for some  $1 \leq t \leq n - 2$ , we obtain

$$\begin{aligned} x_n &\leq A_t x_{n-t} - B_t \sum_{i=0}^{n-t-1} x_i \\ &\leq A_t (r x_{n-t-1} - \sum_{i=0}^{n-t-1} x_i) - B_t \sum_{i=0}^{n-t-1} x_i \\ &= (A_t (r - 1) - B_t) x_{n-t-1} - (A_t + B_t) \sum_{i=0}^{n-t-2} x_i \\ &= A_{t+1} x_{n-t-1} - B_{t+1} \sum_{i=0}^{n-t-2} x_i \end{aligned}$$

320  
321 and similarly for  $y_n$ , thereby establishing (6) for all  $1 \leq t < n$ .

322 Now, observe that setting  $t = n - 1$  in (6) yields

$$x_n \leq A_{n-1} x_1 - B_{n-1} x_0 \leq A_{n-1} y_1 - B_{n-1} y_0 = y_n, \quad (7)$$

323 for all  $n \geq 2$ , which completes the proof.

324

325 Thus, there only remains to justify the bound (5), which will easily follow  
326 from an explicit representation formula for  $(A_t, B_t)$  based on an eigenvector  
327 decomposition of  $M$ . More precisely, straightforward calculations establish  
328 that the eigenvalues of  $M$  are the two roots  $\zeta_{1,r} \leq \zeta_{2,r}$  of the characteristic  
329 polynomial  $p(\zeta) = \zeta^2 - r\zeta + r$ , as defined also in Definition 1. These roots  
330 are both positive if  $r \geq 4$  and distinct whenever  $r > 4$ . In fact, it is readily  
331 seen that  $\zeta_{2,r} \geq \zeta_{1,r} > 1$ . Moreover, it holds that  $\zeta_{2,r} + \zeta_{1,r} = r$ ,  $\zeta_{2,r}\zeta_{1,r} = r$   
332 and  $(\zeta_{2,r} - 1)(\zeta_{1,r} - 1) = 1$ .

333 When  $r > 4$ , we obtain the eigenvector decomposition

$$334 \begin{pmatrix} A_t \\ B_t \end{pmatrix} = \frac{\zeta_{2,r}^t}{\zeta_{2,r} - \zeta_{1,r}} \begin{pmatrix} \zeta_{2,r} - 1 \\ 1 \end{pmatrix} - \frac{\zeta_{1,r}^t}{\zeta_{2,r} - \zeta_{1,r}} \begin{pmatrix} \zeta_{1,r} - 1 \\ 1 \end{pmatrix} \quad (8)$$

335 for every  $t \geq 1$ , which implies (5) because  $\zeta_{2,r} \geq \zeta_{1,r} > 1$ . Finally, further  
336 letting  $r \rightarrow 4$  yields the representation<sup>2</sup>.

$$\begin{pmatrix} A_t \\ B_t \end{pmatrix} = \begin{pmatrix} 2^t + t2^{t-1} \\ t2^{t-1} \end{pmatrix} \quad (9)$$

337

338 in the case  $r = 4$ , which also validates (5) and thus concludes the proof of  
339 the proposition.  $\square$

340 **Observation 7.** *Observe from (7), (8) and (9) in the preceding proof that*  
341 *one has the convenient representation formulas, for every  $n \geq 2$ ,*

---

<sup>2</sup>We include in the analysis the case  $r = 4$  so as to demonstrate that the approach applies to all robustness values.

$$\begin{aligned}
y_n &= \frac{\zeta_{2,r}^{n-1}(\zeta_{2,r} - 1) - \zeta_{1,r}^{n-1}(\zeta_{1,r} - 1)}{\zeta_{2,r} - \zeta_{1,r}} y_1 - \frac{\zeta_{2,r}^{n-1} - \zeta_{1,r}^{n-1}}{\zeta_{2,r} - \zeta_{1,r}} y_0 \\
&= \zeta_{2,r}^{n-1} \frac{(\zeta_{2,r} - 1)y_1 - y_0}{\zeta_{2,r} - \zeta_{1,r}} - \zeta_{1,r}^{n-1} \frac{(\zeta_{1,r} - 1)y_1 - y_0}{\zeta_{2,r} - \zeta_{1,r}},
\end{aligned} \tag{10}$$

342 if  $r > 4$ , and

$$\begin{aligned}
y_n &= (2^{n-1} + (n-1)2^{n-2})y_1 - (n-1)2^{n-2}y_0 \\
&= 2^{n-1}y_1 + (n-1)2^{n-2}(y_1 - y_0),
\end{aligned} \tag{11}$$

343 when  $r = 4$ . If one further requires that (4) hold for  $n = 1$ , whereby  $y_1 =$   
344  $(r-1)y_0$ , then one finds that

$$y_n = \frac{\zeta_{2,r}^n(\zeta_{2,r} - 1) - \zeta_{1,r}^n(\zeta_{1,r} - 1)}{\zeta_{2,r} - \zeta_{1,r}} y_0,$$

345 if  $r > 4$ , and

$$y_n = (3 \cdot 2^{n-1} + (n-1)2^{n-1}) y_0,$$

346 when  $r = 4$ .

347 We are now ready to formally prove Theorem 5.

348 *Proof of Theorem 5.* First, observe that if  $X = (x_i)_{i=0}^\infty$  is  $r$ -robust, then  
349 from (3) it follows that  $\sum_{i=0}^n x_i \leq r x_{n-1}$ . The upper bounds then follow  
350 directly from Proposition 6, with  $x_0 = y_0 = z_0$  and  $x_1 = y_1 = z_1$ , and the  
351 representation formulas (10) and (11).

352 The lower bounds, in contrast, are more subtle. In order to establish their  
353 validity, we define, for any given integer  $j \geq 0$ , an auxiliary sequence  $(x_i)_{i=0}^\infty$   
354 by

$$x_0 = \sum_{k=0}^j z_k \quad \text{and} \quad x_i = z_{j+i}, \text{ if } i \geq 1.$$

355 In particular, it holds that  $\sum_{i=0}^n x_i \leq r x_{n-1}$ , for all  $n \geq 2$ . Therefore, by  
356 Proposition 6 combined with the formulas for  $y_n$  from Observation 7 we

357 deduce that, for all  $n \geq 2$ ,

$$x_n \leq \zeta_{2,r}^{n-1} \frac{(\zeta_{2,r} - 1)x_1 - x_0}{\zeta_{2,r} - \zeta_{1,r}} - \zeta_{1,r}^{n-1} \frac{(\zeta_{1,r} - 1)x_1 - x_0}{\zeta_{2,r} - \zeta_{1,r}},$$

358 if  $r > 4$ , and  $x_n \leq 2^{n-1}x_1 + (n-1)2^{n-2}(x_1 - x_0)$ , for  $r = 4$ . Since  $(x_i)_{i=0}^\infty$   
 359 is also nonnegative and, as  $n \rightarrow \infty$ , the dominant terms above are  $\zeta_{2,r}^{n-1}$  and  
 360  $(n-1)2^{n-2}$ , we conclude that necessarily  $x_0 \leq (\zeta_{2,r} - 1)x_1$ , for all values  $r \geq 4$ .  
 361 In terms of the original sequence  $(z_i)_{i=0}^\infty$ , observing that  $(\zeta_{2,r} - 1)(\zeta_{1,r} - 1) = 1$ ,  
 362 this yields that

$$(\zeta_{1,r} - 1) \sum_{k=0}^j z_k \leq z_{j+1},$$

363 for every  $j \geq 0$ . In particular, if  $z_{i+1} \geq (\zeta_{1,r} - 1)\zeta_{1,r}^i z_0$  holds for every  
 364  $0 \leq i \leq j$ , then

$$\begin{aligned} z_{j+2} &\geq (\zeta_{1,r} - 1) \sum_{k=0}^{j+1} z_k \geq (\zeta_{1,r} - 1)z_0 + (\zeta_{1,r} - 1)^2 z_0 \sum_{k=1}^{j+1} \zeta_{1,r}^{k-1} \\ &= (\zeta_{1,r} - 1)z_0 + (\zeta_{1,r} - 1)^2 z_0 \frac{\zeta_{1,r}^{j+1} - 1}{\zeta_{1,r} - 1} = (\zeta_{1,r} - 1)\zeta_{1,r}^{j+1} z_0, \end{aligned}$$

367 thereby completing the proof of the lower bounds by induction.  $\square$

368 **Observation 8.** *The bounds of Theorem 5 are asymptotically tight. This is*  
 369 *because Property 2 states that any geometric schedule  $G_b$  with  $b \in [\zeta_{1,r}, \zeta_{2,r}]$*   
 370 *is  $r$ -robust.*

371 Using the calculus of upper limits, we prove a property of merged se-  
 372 quences which allows us to bound the range of  $\alpha_{\bar{x}}$ . The technical proof is  
 373 given in Appendix Appendix B.

374 **Lemma 9.** *Let  $(x_{1,i})_{i=1}^\infty, (x_{2,i})_{i=1}^\infty, \dots, (x_{N,i})_{i=1}^\infty$  be  $N$  positive nondecreasing*  
 375 *sequences satisfying the bounds  $c_A A^i \leq x_{j,i} \leq c_B i B^i$ , for all  $i$  and  $j$ , where*

376  $c_A, c_B > 0$  and  $A, B > 1$  are given constants. Let  $(y_i)_{i=1}^{\infty}$  be the sequence  
 377 obtained by merging all  $(x_{j,i})_{i=1}^{\infty}$  and sorting the resulting set of values in  
 378 nondecreasing order. Then,

$$A^{\frac{1}{N}} \leq \liminf_{i \rightarrow \infty} y_i^{\frac{1}{i}} \leq \limsup_{i \rightarrow \infty} y_i^{\frac{1}{i}} \leq B^{\frac{1}{N}}.$$

379

380 Theorem 5, Lemma 9, and the definition of  $\alpha_{\bar{\mathcal{X}}}$  yield:

381 **Corollary 10.** *Let  $X$  be a  $p$ -processor schedule defined by the set  $\mathcal{X} =$   
 382  $\{X_0, X_1, \dots, X_{p-1}\}$ , where each  $X_j$  is an  $r$ -robust strategy, for a given  $r \geq 4$ .  
 383 Then  $\alpha_{\bar{\mathcal{X}}} \in [\zeta_{1,r}^{1/p}, \zeta_{2,r}^{1/p}]$ .*

384 We now state the main result of this section. Its proof formalizes the  
 385 intuition given at the beginning of the section.

386 **Theorem 11 (Lower Bound).** *Any  $r$ -robust contract schedule with  $k$  queries  
 387 has consistency at least  $c(k, r)$ , where*

$$c(k, r) = \min_x \frac{x^{2^k+1}}{x^{2^k} - 1} \quad \text{subject to} \quad \zeta_{1,r}^{1/2^k} \leq x \leq \zeta_{2,r}^{1/2^k}. \quad (12)$$

388

389 *Proof.* Any schedule  $X$  with  $k$  queries will choose a schedule among a set of at  
 390 most  $2^k$  schedules, say  $\mathcal{X} = \{X_0, \dots, X_{2^k-1}\}$ . For  $X$  to be  $r$ -robust, it must  
 391 be that each  $X_i$ , with  $i \in [0, 2^k - 1]$  is likewise  $r$ -robust; otherwise, maliciously  
 392 generated responses would lead to choosing a schedule of robustness greater  
 393 than  $r$ . Note that the consistency of  $X$  is identical to the acceleration ratio  
 394 of the  $2^k$ -processor schedule that is defined by  $\mathcal{X}$ . Let  $X'$  denote this multi-  
 395 processor schedule. From Lemma 4, we have that  $\text{acc}(X') \geq \frac{\alpha_{\bar{\mathcal{X}}}^{2^k+1}}{\alpha_{\bar{\mathcal{X}}}^{2^k} - 1}$ . Last,

396 since every single-processor schedule  $X_i$  must be  $r$ -robust, from Corollary 10  
 397 it follows that  $\alpha_{\bar{\mathcal{X}}} \in [\zeta_{1,r}^{1/2^k}, \zeta_{2,r}^{1/2^k}]$ .  $\square$

398 We will argue that the lower bound of Theorem 11 matches the upper  
 399 bound of [19]. Specifically, consider the following class of single-processor  
 400 schedules.

401 **Definition 12.** For given  $b > 1$ , and  $l \in \mathbb{N}^+$  define  $\mathcal{X}_{b,l}$  as the set of schedules  
 402  $\{X_0, \dots, X_{l-1}\}$ , in which  $X_i = (b^{i+j^l})_{j \geq 0}$ , for every  $i \in [0, l-1]$ .

403 In other words, each schedule  $X_i$  in  $\mathcal{X}_{b,l}$  is a geometric schedule with base  
 404  $b^l$ , whose lengths are multiplied by a factor  $b^i$ . In [19], a  $b > 1$  is chosen so  
 405 as to guarantee that all strategies in  $\mathcal{X}_{b,2^k}$  are  $r$ -robust, and if the response  
 406 string is error-free, then the acceleration ratio of  $\mathcal{X}_{b,2^k}$  is at most  $\frac{b^{2^k+1}}{b^{2^k}-1}$ . Thus,  
 407 there is an  $r$ -robust schedule of consistency at most

$$\min_b \frac{b^{2^k+1}}{b^{2^k}-1} \quad \text{subject to} \quad \zeta_{1,r}^{1/2^k} \leq b \leq \zeta_{2,r}^{1/2^k} \quad \text{and} \quad b > 1. \quad (13)$$

408 The tightness of our lower bound follows from directly comparing (13) with  
 409 Theorem 11 and the fact that  $\zeta_{1,r} > 1$ . We can also express the optimal  
 410 consistency given by (13), and thus obtain the following result.

411 **Theorem 13** (Pareto-optimality). *The optimal consistency of an  $r$ -robust*  
 412 *schedule with  $k$  queries is equal to*

$$c(k, r) = \frac{b^{2^k+1}}{b^{2^k}-1}, \quad \text{with } b = \begin{cases} \zeta_{2,r}^{1/2^k}, & \text{if } r \leq \frac{(1+2^k)^2}{2^k} \\ (1+2^k)^{1/2^k}, & \text{otherwise.} \end{cases} \quad (14)$$

413

414 Figure 2 depicts the optimal consistency-robustness tradeoff as expressed  
 415 by (14) for various values of  $k$ .




Figure 2: Illustration of the optimal consistency-robustness tradeoff.

416 **Example 14.** For  $r = 4$ , the optimal consistency is  $2^{1+\frac{1}{2^k}}$ , thus rapidly  
 417 converging to 2, as function of  $k$ . For large values of  $r$ , the consistency  
 418 converges to  $(1 + 2^k)^{1+1/2^k} / 2^k$ . E.g., for  $k = 4$ , the consistency converges to  
 419 1.26833, as  $r \rightarrow \infty$ .

420 Last, we show that the same Pareto-optimal trade-off can be obtained for  
 421 online bidding via a reduction from contract scheduling.

422 **Theorem 15.** Any  $r$ -robust online bidding strategy, with  $k$  queries, has con-  
 423 sistency at least  $c(k, r)$ , as expressed by (14). Furthermore, this result is tight.

424

425 *Proof.* For the lower bound, we will show a reduction from contract schedul-  
 426 ing with  $k$  queries. Let  $T$  be the unknown interruption time for a given  
 427 contract scheduling instance. Let  $\mathcal{X}$  be an  $r$ -robust,  $c$ -consistent bidding  
 428 strategy with  $k$  queries, for some given  $r, c$ . We will show how to obtain a  
 429 schedule with the same guarantees. To this end, we can interpret  $\mathcal{X}$  as a set  
 430 of  $2^k$  sequences  $\{X_0, \dots, X_{2^k-1}\}$ , each of which must be  $r$ -robust. We define  
 431 the bidding target to be equal to  $T$ , and let  $Y \in \mathcal{X}$  be the sequence returned



432 by the query responses for this bidding instance  $u = T$ .

433 By definition,  $Y$  is  $r$ -robust. Let  $m$  denote the smallest index for which  
 434  $y_m \geq T$ . Since  $\mathcal{X}$  is also  $c$ -consistent, we have  $\sum_{j=0}^i y_j \leq r y_{i-1}$ , for all  $i$ , and  
 435  $\sum_{j=0}^m y_j \leq cT$ . Define  $Z = (z_i)_{i \geq 0}$  such that  $z_i = y_i/c$ . It follows that

$$\sum_{j=0}^i z_j \leq r z_{i-1}, \text{ for all } i, \text{ and } z_m \geq \frac{T}{c}, \quad \sum_{j=0}^m z_j \leq T. \quad (15)$$

436 If we interpret  $Z$  as a contract schedule, the first equation in (15) shows that  
 437  $Z$  is  $r$ -robust, whereas the last two equations show that it is  $c$ -consistent.  
 438 Therefore, from Theorem 11, it follows that  $c \geq c(k, r)$ .  $\square$

439 Last, we note that [14] gave an explicit strategy for online bidding of  
 440 consistency equal to  $c(k, r)$ , which establishes the tightness of the result.

#### 441 4. Sequencing with noisy queries

442 In this section, we extend our study to the noisy query model. We focus  
 443 on contract scheduling, but we note that the same upper bounds can be  
 444 easily extended to online bidding using the same techniques. Recall that  
 445 in this model,  $\eta \leq k$  query responses may be erroneous, and the scheduler  
 446 specifies a parameter  $\tau \leq k$  that describes its desired tolerance to errors or,  
 447 alternatively, an anticipated upper bound on the query error.

448 We first discuss the intuition behind our approach. The starting obser-  
 449 vation is that the Pareto-optimal schedule can be found in the class  $\mathcal{X}_{b,2^k}$   
 450 (recall Definition 12), as shown in Section 3). This class has a nice structural  
 451 property, as we show in Property 16: if  $j_T$  is the index of the best schedule  
 452 in this class, for interruption  $T$ , then any schedule of index close, but smaller

453 than  $j_T$  is likewise quite efficient. In the next step, we show how to apply  
 454 an *error-tolerant* binary-search algorithm, using  $k$  binary queries in the pres-  
 455 ence of errors, so as to find a schedule in  $\mathcal{X}_{b,2^k}$  that has index close to (but no  
 456 larger than)  $j_T$ ; Property 16 implies that this schedule should also perform  
 457 well. In Theorem 18, we quantify this statement and optimize the choice of  
 458  $b$ . Our main result is Theorem 20: here, we prove that the obtained schedule  
 459 is extremely robust to errors and very close to the Pareto-optimal one as  $k$   
 460 increases. It is worth noting that the resulting schedule is determined by  
 461 *adaptive* queries since the underlying binary search algorithm builds upon  
 462 adaptive queries. Adaptivity, in particular, helps us find an efficient schedule  
 463 within the exponential-sized class  $\mathcal{X}_{b,2^k}$ .

464 We start with the structural property that follows immediately from Def-  
 465 inition 12. See also Figure 3 for an illustration.

466 **Property 16.** *Consider the set  $\mathcal{X}_{b,l} = \{X_0, \dots, X_{l-1}\}$  of the  $l$  single-processor  
 467 strategies in Definition 12. For a given interruption time  $T$ , let  $j_T \in [0, l-1]$   
 468 be such that  $X_{j_T}$  completes a contract of largest length among all schedules  
 469 in  $\mathcal{X}_{b,l}$ , and specifically of length  $L_0$ . Then, for all  $j \in [0, j_T]$ ,  $X_j$  completes  
 470 a contract of length at least  $b^{j-j_T} L_0$  by  $T$ .*

471 We emphasize that  $j_T$  depends on the interruption time  $T$ , and is thus  
 472 unknown to the scheduler. We would like to find a schedule of index close, but  
 473 no larger than  $j_T$ . To model this situation, we define a problem which we call  
 474 MINSEARCH. The input to this problem is an unknown array  $A[0 \dots 2^k - 1]$   
 475 that stores a permutation of  $\{0, \dots, 2^k - 1\}$ , such that there exists a  $j^* \in$   
 476  $[0, 2^k - 1]$  for which  $A[j^* - i] = i$ , for all  $i \in [0, j^*]$ . For example,  $A$  can  
 477 be of the form  $A = [3, 2, 1, 0, 6, 4, 7, 5]$ , where  $j^* = 3$ . The objective is to




Figure 3: An illustration of the class of schedules  $\mathcal{X}_{b,2^k}$ , for  $k = 2$ . Given an interruption  $T$ , the schedule that completes the largest contract by time  $T$  is  $X_1$  (a contract of length  $b^9$ ) hence  $j_T = 1$ .

478 use  $k$  binary queries, at most  $\tau$  of which can receive erroneous responses,  
 479 for a given  $\tau$ , so as to identify an index  $i \in [0, j^*]$  such that  $A[i]$  is as small  
 480 as possible, without knowing  $j^*$ . In [40], an algorithm called *Robust Binary*  
 481 *Interval Search* (RBIS) was given in the context of the online *time-series*  
 482 *search* problem, based on an algorithm for a related error-tolerant setting  
 483 due to [38]. The same algorithm can be applied to MINSEARCH:

484 **Observation 17.** *There is an algorithm for MINSEARCH with  $k$  binary*  
 485 *queries that outputs an index  $i \in [0, j^*]$  such that  $A[i] \leq 2^{\lfloor k/2 \rfloor + 2\tau}$ , for all*  
 486  *$\tau \leq k/4$ , if there are at most  $\tau$  erroneous query responses.*

487 Thus, in our schedule, each binary query is of the form “is  $j_T \leq x$ ?”,  
 488 where  $x$  is chosen in  $[0, 2^k - 1]$ . Note, however, that from the structure of  
 489  $\mathcal{X}_{b,2^k}$ , one can express equivalently such query as a partition query of the  
 490 form “is the interruption in  $A$  or  $B$ ?”, where  $A, B$  form a disjoint partition of  
 491 the timeline. We refer to [19] for the details on the implementation of such  
 492 partition queries.

493 Combining Property 16 and Remark 17, we obtain the following perfor-

494 mance guarantees. Define  $U = 2^{\lfloor k/2 \rfloor + 2\tau}$ .

495 **Theorem 18.** *For any  $r \geq 4$  and  $\eta \leq \tau \leq k/4$ , there exists a schedule that*  
 496 *has acceleration ratio at most*

$$R(k, r, \tau) = \frac{b^{2^k + U + 1}}{b^{2^k} - 1}, \quad (16)$$

497 where

$$b = \begin{cases} \zeta_{2,r}^{1/2^k}, & \text{if } r \leq \frac{(1+2^k/(1+U))^2}{2^k/(1+U)} \\ (1 + 2^k/(1+U))^{1/2^k}, & \text{otherwise,} \end{cases}$$

498 and where  $U = 2^{\lfloor k/2 \rfloor + 2\tau}$ . Otherwise, the acceleration ratio of the schedule is  
 499 at most  $r$ .

500 *Proof.* We apply the error-tolerant binary search algorithm of Remark 17 for  
 501 MINSEARCH on the set of indices of all schedules in  $\mathcal{X}_{b,2^k}$ , where  $b > 1$  will  
 502 be chosen later. Assuming that  $\eta \leq \tau$ , then the output is the index of a  
 503 strategy in  $\mathcal{X}_{b,2^k}$  which is ranked at most  $U$  among the schedules in  $\mathcal{X}_{b,2^k}$ ,  
 504 in regards to its largest completed contract (where lower ranking indicates  
 505 better performance). From Property 16, this means that the selected schedule  
 506 has length at least  $L_0/b^U$ , where  $L_0$  is the length of the largest contract  
 507 completed by time  $T$  among all schedules in  $\mathcal{X}_{b,2^k}$ . It is easy to evaluate the  
 508 latter formally (see the details in the proof of Theorem 22), and we infer that  
 509 the acceleration ratio of the selected schedule is at most  $\frac{b^{2^k + 1 + U}}{b^{2^k} - 1}$ .

510 Furthermore, given the desired robustness  $r \geq 4$ , we require that each  
 511 schedule in  $\mathcal{X}_{b,2^k}$  must be  $r$ -robust. From Property 2, this is equivalent to  
 512 the constraint  $\zeta_{1,r} \leq b^{2^k} \leq \zeta_{2,r}$ . Therefore, the acceleration ratio of the  
 513 schedule is equal to

$$\min \frac{b^{2^k + 1 + U}}{b^{2^k} - 1}, \text{ subject to } \zeta_{1,r} \leq b^{2^k} \leq \zeta_{2,r}.$$

514 Optimizing the above expression yields the result. □

515 The schedule of Theorem 18 is very robust to query errors. More precisely,  
 516 in Theorem 20, we show that as long as  $\eta \leq \tau \leq k/4 - o(k)$ , not only does the  
 517 acceleration ratio degrade gently as a function of  $\tau$ , but, more importantly, it  
 518 improves rapidly as  $k$  increases, and approaches the ideal performance of the  
 519 Pareto-optimal schedule. Recall that  $c(k, r)$  and  $R(k, r, \tau)$  are given by (14)  
 520 and (16), respectively. We first define a function that will help us express  
 521 and prove the performance guarantee.

522 **Definition 19.** For all  $x \geq 1$ , define the function  $g(x) = \frac{(1+x)^{1+\frac{1}{x}}}{x}$ .

523 Note that  $g$  is a decreasing function of  $x$  and converges to 1 as  $x \rightarrow \infty$ .

524 **Theorem 20.** For all  $k \in \mathbb{N}^+$ ,  $r \geq 4$  and  $\eta \leq \tau \leq k/4$ , it holds that

$$R(k, r, \tau) \leq f(k, r, \tau) \cdot c(k, r),$$

525 where  $f(k, r, \tau) = \frac{g(\frac{2^k}{1+U})}{g(2^k)} \cdot \zeta_{2,r}^{\frac{U}{2^k}}$ .

526 *Proof.* We consider cases, depending on whether  $r \leq \frac{(1+2^k)^2}{2^k}$  or not (see (14))  
 527 and whether  $r \leq \frac{(1+2^k/(1+U))^2}{2^k/(1+U)}$  or not (see (16)). To simplify notation, define

$$\rho_0 = \frac{(1+2^k)^2}{2^k} \quad \text{and} \quad \rho_1 = \frac{(1+2^k/(1+U))^2}{2^k/(1+U)},$$

528 and note that it is always the case  $\rho_0 \geq \rho_1$ , since the function  $(1+x^2)/x$   
 529 is increasing for all  $x \geq 1$ . In the discussion that follows, recall that the  
 530 function is defined in Definition 19.

531 *Case 1:*  $r \leq \rho_1$  (thus  $r \leq \rho_0$  as well). In this case,

$$R(k, r, \tau) = \frac{\zeta_{2,r}^{1+\frac{1+U}{2^k}}}{\zeta_{2,r} - 1}, \quad \text{and} \quad c(k, r) = \frac{\zeta_{2,r}^{1+\frac{1}{2^k}}}{\zeta_{2,r} - 1}.$$

532 Therefore, we have that  $R(k, r, \tau) = \zeta_{2,r}^{\frac{U}{2^k}} c(k, r)$ , and the theorem follows .

533 *Case 2:*  $r > \rho_0$  (and thus  $r > \rho_1$  as well). In this case,

$$R(k, r, \tau) = g(2^k/(1+U)), \quad \text{and} \quad c(k, r) = g(2^k),$$

534 therefore we have that

$$\frac{R(k, r, \tau)}{c(k, r)} = \frac{g(\frac{2^k}{1+U})}{g(2^k)},$$

535 and the theorem follows.

536 *Case 3:*  $r \in (\rho_1, \rho_0]$ . For this case to apply, it must be that

$$\zeta_{2,r} - 1 \leq 2^k \quad \text{and} \quad \zeta_{2,r} - 1 \geq \frac{2^k}{1+U}. \quad (17)$$

537 We have that

$$R(k, r, \tau) = g\left(\frac{2^k}{1+U}\right) \quad \text{and} \quad c(k, r, \tau) = \frac{\zeta_{2,r}^{1+\frac{1}{2^k}}}{\zeta_{2,r} - 1}, \quad (18)$$

538 Define  $y = \frac{2^k}{1+U}$  and  $z = \zeta_{2,r} - 1$ , then by (17) we have that  $z \geq y$ . Further-  
539 more, we can express  $c(k, r)$  equivalently as

$$c(k, r, \tau) = g(z) \frac{\zeta_{2,r}^{\frac{1}{2^k}}}{\zeta_{2,r}^{\frac{1}{z}}}. \quad (19)$$

540 Combining the above, we obtain

$$\begin{aligned} \frac{R(k, r, \tau)}{c(k, r)} &= \frac{g(y)}{g(z)} \cdot \zeta_{2,r}^{\frac{1}{z} - \frac{1}{2^k}} && \text{(From (18) and (19))} \\ &= \frac{g(y)}{g(z)} \cdot \zeta_{2,r}^{\frac{U}{2^k}} && \text{(From (17))} \\ &= \frac{g(\frac{2^k}{1+U})}{g(\zeta_{2,r} - 1)} \cdot \zeta_{2,r}^{\frac{U}{2^k}} && \text{(def of } z, y) \\ &\leq \frac{g(\frac{2^k}{1+U})}{g(2^k)} \cdot \zeta_{2,r}^{\frac{U}{2^k}}. && \text{(From monotonicity of } g \text{ and (17))} \end{aligned}$$

541

□

542 Theorem 20 leads to interesting conclusions in regard to the robustness of  
 543 the schedule to errors. Specifically, simple calculus shows that the function  
 544 rapidly converges to 1, as  $k \rightarrow \infty$ , even if  $\tau$  is as large as  $k/4 - o(k)$ . This  
 545 implies that the scheduler is not constrained by the choice of a particular  
 546 tolerance parameter  $\tau$ , even if  $k$  is a fairly small constant. Namely, even if  
 547 we choose a very pessimistic value for  $\tau$  such as  $\tau = k/4 - o(k)$ , and even if  
 548 the error is as high as  $k/4 - o(k)$ , the schedule has an acceleration ratio very  
 549 close to the ideal case of error-free queries (i.e., as good as its consistency).  
 550 This improves considerably upon the state-of-the-art schedule:

551 **Observation 21.** *The schedule of [19] does not exhibit comparable robust-*  
 552 *ness. To see this, consider the case  $r = 4$ , for which their acceleration ratio*  
 553 *is  $2^{1+\frac{1}{k}+2\frac{\tau}{k}}$ . Suppose that  $\tau = \Theta(k)$ , i.e.,  $\tau = c \cdot k$  for constant  $c < 1/4$ , then*  
 554 *this is equal to  $2^{1+\frac{1}{k}+2c}$ , which not only can be much larger than the ideal*  
 555 *consistency, if  $c$  is large, but is also practically unaffected by the number of*  
 556 *queries  $k$ .*

## 557 5. Robust fault-tolerant contract scheduling

558 The techniques we introduced in Sections 3 and 4 can be applied to other  
 559 multi-criteria optimization settings unrelated to query predictions. We il-  
 560 lustrate, using as an example the *robust, fault-tolerant contract scheduling*  
 561 *problem*, defined as follows. Suppose we are given a system of  $p$  identical  
 562 parallel processors, a *robustness* requirement  $r$ , and a *fault-tolerance* param-  
 563 eter  $f < p$ . The objective is to obtain a schedule of contract algorithms in  
 564 the system of  $p$  parallel processors that has a minimum acceleration ratio if  
 565 up to  $f$  processors can be faulty but also has an acceleration ratio at most

566  $r$  if all but a single processor are faulty. Here, a fault is byzantine, in that a  
 567 processor may cease to work at any point in time. We denote this problem  
 568 as  $\text{RFT}(r, p, f)$ . We emphasize that there is no concept of predictions in this  
 569 formulation.

570 This problem generalizes the fault-tolerant setting studied by Kupavskii  
 571 and Welzl [37]; in their model, there is no robustness requirement, and one  
 572 aims to minimize the acceleration ratio of the system in the presence of at  
 573 most  $f$  faults. In our setting, instead, we treat  $f$  as a “soft” bound on the  
 574 number of faults that may occur and would still like the schedule to perform  
 575 well if this bound is exceeded. Note that if  $r = \infty$ ,  $\text{RFT}$  reduces to the setting  
 576 of [37].

577 We first show how to obtain a schedule for  $\text{RFT}(r, p, f)$  by analyzing an  
 578 explicit schedule. Namely, for a  $b > 1$  to be fixed later, we will use the family  
 579 of schedules  $\mathcal{X}_{b,p}$  (recall Definition 12), with each strategy  $X_i$  in this collection  
 580 scheduled on processor  $i$ . From Property 16, the acceleration ratio of this  
 581 schedule is  $b^{p+f+1}/(b^p - 1)$ , in the presence of at most  $f$  faults. Moreover,  
 582 since each  $X_i$  is near-geometric, we can enforce the requirement that its  
 583 individual acceleration ratio is at most  $r$  by appealing to Property 2. This  
 584 yields the following result.

585 **Theorem 22.** *There is a schedule for  $\text{RFT}(r, p, f)$  of acceleration ratio at*  
 586 *most*

$$\min \frac{b^{p+f+1}}{b^p - 1} \quad \text{subject to} \quad b \in [\zeta_{1,r}^{1/p}, \zeta_{2,r}^{1/p}].$$

587

588 *Proof.* Consider an interruption  $T$  that occurs right before the contract of  
 589 length  $b^{j+p+l}$  terminates, for some  $j \in \mathbb{N}$ , and  $l \in [0, p - 1]$ . Then, from the



590 definition of  $\mathcal{X}_{b,p}$ , a contract of length  $b^{jp+l-f-1}$  has completed, even if  $f$   
 591 processor faults have occurred. We have

$$\frac{T}{b^{jp+l-f-1}} = \frac{\sum_{i=0}^j b^{ip+l}}{b^{jp+l-f-1}} \leq \frac{b^{p+f+1}}{b^p - 1}.$$

592 which, together with Property 2 completes the proof.  $\square$

593 We will show that this result is tight, and thus, our schedule is optimal.  
 594 To this end, we need the following generalization of Lemma 4, to the fault-  
 595 tolerant setting. The proof closely follows an approach from [41], which  
 596 studied a fault-tolerant version of the online bidding problem. The technical  
 597 proof is given in Appendix Appendix B for completeness.

598 **Lemma 23.** [Appendix] *Let  $X$  be a  $p$ -processor schedule, as defined by a*  
 599 *set  $\mathcal{X}$  of  $p$  single-processor strategies, assuming that each strategy in  $\mathcal{X}$  has*  
 600 *finite acceleration ratio and that at most  $f$  processors may be faulty. Then*  
 601  $\text{acc}(X) \geq \frac{\alpha_{\mathcal{X}}^{p+f+1}}{\alpha_{\mathcal{X}}^p - 1}.$

602 We can now show that the upper bound of Theorem 22 is tight.

603 **Theorem 24.** *Any schedule for  $\text{RFT}(r, p, f)$  of acceleration ratio at most*

$$\min \frac{b^{p+f+1}}{b^p - 1} \quad \text{subject to} \quad b \in [\zeta_{1,r}^{1/p}, \zeta_{2,r}^{1/p}].$$

604

605 *Proof.* Given a  $p$ -processor schedule  $X$ , and from the robustness requirement  
 606 of the problem, the schedule of each individual processor must be  $r$ -robust.  
 607 Thus, Corollary 10 applies, and in combination with Lemma 23, we obtain  
 608 that the acceleration ratio of any schedule is at least  $\min \frac{\alpha_{\mathcal{X}}^{p+f+1}}{\alpha_{\mathcal{X}}^p - 1}$ , subject to  
 609 the condition  $\alpha_{\mathcal{X}} \in [\zeta_{1,r}^{1/p}, \zeta_{2,r}^{1/p}]$ , which completes the proof.  $\square$

610 We can obtain a closed-form expression of the acceleration ratio of The-  
611 orem 22 along the lines of the result of Theorem 18, by replacing  $2^k$  with  $p$ ,  
612 and  $U$  with  $f$ .

613 **Example 25.** Consider  $\text{RFT}(r, p, f)$  with  $r = 4$ ,  $p = 3$  and  $f = 1$ . Then  
614 Theorem 22 shows that there is a 3-processor schedule which has acceleration  
615 ratio  $2^{5/3} \approx 3.175$ , if at most one processor fault occurs, and has acceleration  
616 ratio 4 if all but one processor may be faulty.

### 617 5.1. A note on fault-tolerant contract scheduling

618 We emphasize that our results provide rigorous proofs not only for RFT,  
619 but also for the standard fault-tolerant contract scheduling problem intro-  
620 duced in [37] (recall that in the latter, we have  $r = \infty$ ). We note that [37] did  
621 not provide explicit acceleration ratios but gave an optimal strategy for the  
622 problem of searching for a hidden target in the line in a setting in which there  
623 are  $p$  searchers,  $f$  of which may be faulty. The work implies that the same ap-  
624 proach gives a solution to the problem of fault-tolerant contract scheduling,  
625 however we argue below that the two problems are fundamentally differ-  
626 ent in multi-processor/multi-searcher settings, and it is not obvious how to  
627 reduce one to the other.

628 Consider first the case  $f = 0$ , i.e., no searchers or processors are faulty.  
629 Then, with only two searchers, the competitive ratio of searching on the line  
630 is 1: one can dedicate a searcher to each direction of the line. In contrast, [6]  
631 shows that in any  $p$ -processor contract schedule, the acceleration ratio is  
632 always strictly larger than one. It is also instructive to see that the proof of [6]  
633 is very much different than, say, [42], which studied multi-robot searching.

634 This observation extends to the fault-tolerant setting: while it is possible  
635 to explore each of the two branches of the line in an optimal manner, by  
636 dedicating  $f + 1$  searchers to search that branch exclusively, this action has  
637 no counterpart in the domain of contract scheduling, since infinite length  
638 executions of contract algorithms lead to schedules of unbounded acceleration  
639 ratio. Nevertheless, our results in Section 5 apply not only to robust, fault-  
640 tolerant contract scheduling but also provide explicit formulas for the vanilla  
641 fault-tolerant model of [37].

## 642 6. Experimental evaluation

643 We present an experimental analysis of the query-based schedule we in-  
644 troduced in Section 4; we refer to this schedule as ROBUST QUERY-BASED  
645 SCHEDULE, or RQS for brevity (and recall its analysis in Theorem 18). This  
646 is the most general setting studied in this work and the only one that incor-  
647 porates noisy data and randomness.

648 We evaluate RQS for interruptions  $T \in (1, 10^5]$ . More precisely, we con-  
649 sider all  $T$  of the form  $T = \lceil 1.01^i \rceil$  for  $i \in [1, 1157]$ . For each such potential  
650 interruption, there is a unique, error-free bit string of size  $k$ , which identi-  
651 fies the best schedule among the  $2^k$  candidates in  $\mathcal{X}_{b,2^k}$  ( $b$  chosen according  
652 to (16)), and whose bits are responses to binary partition queries, as ex-  
653 plained in Section 4. Given this string and a tolerance parameter  $\tau \leq k/4$ ,  
654 we generate the noisy response to the  $k$  queries by choosing  $\eta$  uniformly at  
655 random in  $[0, \tau]$ , then flipping  $\eta$  bits, again chosen uniformly at random<sup>3</sup>.

---

<sup>3</sup>Since RQS is determined by a noisy search over a space of  $2^k$  schedules, the experi-  
mental evaluation needs to assume a reasonably small value of  $k$ .




Figure 4: Experimental performance ratios of various schedules for the setting  $r = 4$ ,  $k = 16$ , and  $\tau \in \{2, 4\}$ .

656 For each interruption  $T$  and its associated noisy query responses, we  
657 evaluate the *experimental* performance ratio, namely the ratio between  $T$   
658 and the longest contract completed in the schedule at hand by time  $T$ . More  
659 precisely, for each  $T$ , we generate 100 random noisy query responses (each  
660 a  $k$ -bit string) as described above, and the longest completed contract in  
661 the schedule is the average one over these 100 runs. For a given robustness  
662 guarantee  $r \geq 4$ , we compare RQS against the schedule  $Robust_\tau$  of [19], to  
663 which we refer as *Baseline Robust*. We also compare RQS to two *oblivious*  
664 schedules that do not rely to queries, namely the geometric schedules with  
665 base  $\zeta_{1,r}$  and  $\zeta_{2,r}$ , respectively. Recall that, from Property 2, these are the  
666 two extreme base values for which a geometric schedule without queries is  
667  $r$ -robust, and note that for  $r = 4$ ,  $\zeta_{1,4} = \zeta_{2,4} = 2$  (i.e., there is a single  
668 oblivious schedule).

669 Figure 4 depicts the experimental performance ratio of the various sched-  
670 ules for  $k = 16$ ,  $r = 4$ , and  $\tau \in \{2, 4\}$ . The peaks of the plots are the  
671 empirical acceleration ratios of the corresponding schedules. RQS improves

672 upon Baseline-Robust and Oblivious in 94.76% and 94.94% of interruption  
673 times, respectively, for  $\tau = 2$ , and 95.77% and 94.94%, respectively, for  $\tau = 4$ .  
674 For  $\tau = 2$ , we observe that RQS and Baseline-Robust have experimental ac-  
675 celeration ratios 2.054 and 2.355, whereas for  $\tau = 4$ , their acceleration ratios  
676 are 2.402 and 2.596, respectively.

677 The results demonstrate that for small values of error, i.e., for  $\tau = 2$ , the  
678 empirical performance ratio of RQS is very close to 2, which is consistent  
679 with Theorems 11 and 20, since the ideal consistency (assuming no error)  
680 is equal to  $2^{1+\frac{1}{2k}}$ . As the error increases, i.e., for  $\tau = 4$ , the performance  
681 of both schedules becomes more noisy, as expected; however, RQS still per-  
682 forms better than Baseline-Robust, even for  $\tau = k/4$ . Hence, as long as the  
683 query error is not prohibitive, the improvements over Baseline-Robust are  
684 considerable.

685 We refer to Appendix Appendix C for many additional experimental  
686 results. In particular, we show that the performance improvement becomes  
687 more pronounced as  $r$  increases. For example, for  $r = 6$  and  $r = 10$ , RQS  
688 performs better than Baseline-Robust in at least 98% of interruptions, and  
689 the acceleration ratio improvement is at least 16%. This is due to the fact  
690 that RQS has greater leeway to optimize the base  $b$ , given that the worst-case  
691 effect of the query error is less significant than in Baseline-Robust. We also  
692 show experiments with different numbers of queries, ranges of query errors,  
693 and query errors either below or beyond the tolerance parameter. These  
694 results show that, as expected, RQS performs much better than Baseline-  
695 Robust if  $\eta$  is much smaller than  $\tau$ , and it is comparable to Baseline-  
696 Robust even if  $\eta$  exceeds  $\tau$ .

697 **7. Discussion**

698 We studied a classic problem from sequential decision-making in the  
699 query-based prediction model. Our approach exploits connections between  
700 the quality of responses, parallelism, error-tolerant search, and fault-tolerance  
701 at the processor level. We gave Pareto-optimal schedules, and schedules  
702 based on adaptive queries in which the performance degradation due to error  
703 is negligible, as the number of queries grows. We also showed that our tech-  
704 niques are applicable to multi-criteria scheduling beyond the query setting,  
705 such as the robust, fault-tolerant contract scheduling problem.

706 The techniques of Section 3 can be useful in *best-of-both-worlds* analysis  
707 and multi-criteria optimization for other problems such as searching for a  
708 hidden target under the competitive ratio in an unbounded line environment,  
709 a classic search problem in TCS, AI and OR that has been studied in a variety  
710 of settings. More precisely, we know from [3, 10] that any contract schedule  
711 of acceleration ratio  $r$  can be interpreted as a search strategy of competitive  
712 ratio  $1 + 2r$ , and vice versa. Theorem 5 then solves the following problem: it  
713 characterizes the  $(1 + 2r)$ -competitive strategies whose search lengths are as  
714 large as possible (the upper bound in the statement of the theorem), or as  
715 small as possible (the lower bound, respectively). This answers the question  
716 of finding the most “aggressive” and the most “conservative” strategies for  
717 any desired competitive ratio. Such characterizations help obtain strategies  
718 that simultaneously optimize multiple performance measures; see., e.g., [43],  
719 which answered this question only for the special case  $r = 9$ .

720 The broader future objective is to obtain learning-augmented algorithms  
721 in settings in which predictions are elicited via queries. For instance, [11]

722 showed connections between sequencing problems (such as contract schedul-  
723 ing) and problems such as hierarchical clustering and minimum latency,  
724 which opens the possibility of applying our approaches to other important  
725 problems.

## 726 **References**

- 727 [1] S. Zilberstein, Using anytime algorithms in intelligent systems, *AI*  
728 *Magazine* 17 (1996) 73–83.
- 729 [2] S. J. Russell, S. Zilberstein, Composing real-time systems, in: *Proc.*  
730 *International Joint Conference on Artificial Intelligence (IJCAI)*, 1991,  
731 pp. 212–217.
- 732 [3] D. S. Bernstein, L. Finkelstein, S. Zilberstein, Contract algorithms and  
733 robots on rays: Unifying two scheduling problems, in: *Proc. Inter-*  
734 *national Joint Conference on Artificial Intelligence (IJCAI)*, 2003, pp.  
735 1211–1217.
- 736 [4] D. S. Bernstein, T. J. Perkins, S. Zilberstein, L. Finkelstein, Scheduling  
737 contract algorithms on multiple processors, in: *Proc. AAAI Conference*  
738 *on Artificial Intelligence*, 2002, pp. 702–706.
- 739 [5] S. Zilberstein, F. Charpillet, P. Chassaing, Optimal sequencing of con-  
740 tract algorithms., *Ann. Math. Artif. Intell.* 39 (2003) 1–18.
- 741 [6] A. López-Ortiz, S. Angelopoulos, A. Hamel, Optimal scheduling of con-  
742 tract algorithms for anytime problem-solving, *J. Artif. Intell. Res.* (2014)  
743 533–554.

- 744 [7] S. Angelopoulos, A. López-Ortiz, A. M. Hamel, Optimal scheduling of  
745 contract algorithms with soft deadlines, *J. Sched.* 20 (2017) 267–277.
- 746 [8] S. Angelopoulos, A. López-Ortiz, Interruptible algorithms for multi-  
747 problem solving, *J. Sched.* 23 (2020) 451–464.
- 748 [9] S. Angelopoulos, S. Jin, Earliest-completion scheduling of contract al-  
749 gorithms with end guarantees, in: *Proc. International Joint Conference*  
750 *on Artificial Intelligence, (IJCAI)*, 2019, pp. 5493–5499.
- 751 [10] S. Angelopoulos, Further connections between contract-scheduling and  
752 ray-searching problems, in: *Proc. International Joint Conference on*  
753 *Artificial Intelligence (IJCAI)*, 2015, pp. 1516–1522.
- 754 [11] M. Chrobak, C. Kenyon-Mathieu, *SIGACT news online algorithms col-*  
755 *umn 10: Competitiveness via doubling*, *SIGACT News* 37 (2006) 115–  
756 126.
- 757 [12] M. Chrobak, C. Kenyon, J. Noga, N. E. Young, Incremental medians  
758 via online bidding, *Algorithmica* 50 (2008) 455–478.
- 759 [13] L. Epstein, A. Levin, Randomized algorithms for online bounded bid-  
760 ding, *Information processing letters* 110 (2010) 503–506.
- 761 [14] S. Angelopoulos, C. Dürr, S. Jin, S. Kamali, M. P. Renault, Online  
762 computation with untrusted advice, in: *Proc. International Conference*  
763 *on Innovations in Theoretical Computer Science (ITCS)*, 2020, pp. 52:1–  
764 52:15.



- 765 [15] S. Im, B. Moseley, C. Xu, R. Zhang, Online state exploration: Competi-  
766 tive worst case and learning-augmented algorithms, in: Proc. European  
767 Conference on Machine Learning and Knowledge Discovery in Databases  
768 (ECML/PKDD), 2023, pp. 333–348.
- 769 [16] K. Anand, R. Ge, A. Kumar, D. Panigrahi, A regression approach to  
770 learning-augmented online algorithms, in: Advances in Neural Informa-  
771 tion Processing Systems 34: Annual Conference on Neural Information  
772 Processing Systems 2021, NeurIPS 2021, 2021, pp. 30504–30517.
- 773 [17] A. Beck, D. Newman, Yet more on the linear search problem, Israel  
774 Journal of Mathematics 8 (1970) 419–429.
- 775 [18] M. X. Goemans, J. M. Kleinberg, An improved approximation ratio for  
776 the minimum latency problem, Mathematical Programming 82 (1998)  
777 111–124.
- 778 [19] S. Angelopoulos, S. Kamali, Contract scheduling with predictions, J.  
779 Artif. Intell. Res. 77 (2023) 395–426.
- 780 [20] T. Lykouris, S. Vassilvitskii, Competitive caching with machine learned  
781 advice, J. ACM 68 (2021) 24:1–24:25.
- 782 [21] M. Purohit, Z. Svitkina, R. Kumar, Improving online algorithms via  
783 ML predictions, in: Proc. Conference on Neural Information Processing  
784 Systems (NIPS), 2018, pp. 9661–9670.
- 785 [22] M. Mitzenmacher, S. Vassilvitskii, Algorithms with predictions, in:  
786 Beyond the Worst-Case Analysis of Algorithms, Cambridge University  
787 Press, 2020, pp. 646–662.

- 788 [23] A. Lindermayr, N. Megow, Repository of works on algorithms with pre-  
789 dictions, <https://algorithms-with-predictions.github.io/about>,  
790 2023. Accessed: 2024-04-26.
- 791 [24] A. Mazumdar, B. Saha, Clustering with noisy queries, in: Proc. Con-  
792 ference on Neural Information Processing Systems (NIPS), volume 30,  
793 2017, pp. 5788–5799.
- 794 [25] R. Addanki, S. Galhotra, B. Saha, How to design robust algorithms  
795 using noisy comparison oracle, Proc. VLDB Endow. 14 (2021) 1703–  
796 1716.
- 797 [26] Z. Benomar, V. Perchet, Advice querying under budget constraint for  
798 online algorithms, in: Advances in Neural Information Processing Sys-  
799 tems 36: Annual Conference on Neural Information Processing Systems  
800 2023, NeurIPS 2023, 2023, pp. 75026 – 75047.
- 801 [27] S. Im, R. Kumar, A. Petety, M. Purohit, Parsimonious learning-  
802 augmented caching, in: International Conference on Machine Learn-  
803 ing, ICML, volume 162 of *Proceedings of Machine Learning Research*,  
804 PMLR, 2022, pp. 9588–9601.
- 805 [28] B. Sun, R. Lee, M. Hajiesmaili, A. Wierman, D. Tsang, Pareto-optimal  
806 learning-augmented algorithms for online conversion problems, in: Proc.  
807 Conference on Neural Information Processing Systems (NeurIPS), 2021,  
808 pp. 10339–10350.
- 809 [29] A. Wei, F. Zhang, Optimal robustness-consistency trade-offs for

- 810 learning-augmented online algorithms, *Advances in Neural Information*  
811 *Processing Systems* 33 (2020) 8042–8053.
- 812 [30] T. Li, R. Yang, G. Qu, G. Shi, C. Yu, A. Wierman, S. H. Low, Robust-  
813 ness and consistency in linear quadratic control with untrusted predic-  
814 tions, *Proc. ACM Meas. Anal. Comput. Syst.* 6 (2022) 18:1–18:35.
- 815 [31] R. Lee, J. Maghakian, M. H. Hajiesmaili, J. Li, R. K. Sitaraman, Z. Liu,  
816 Online peak-aware energy scheduling with untrusted advice, in: *Proc.*  
817 *International Conference on Future Energy Systems (e-Energy)*, 2021,  
818 pp. 107–123.
- 819 [32] N. Christianson, J. Shen, A. Wierman, Optimal robustness-consistency  
820 tradeoffs for learning-augmented metrical task systems, in: *Proc. Inter-*  
821 *national Conference on Artificial Intelligence and Statistics (AISTATS)*,  
822 2023, pp. 9377–9399.
- 823 [33] R. Lee, B. Sun, M. Hajiesmaili, J. C. S. Lui, Online search with predic-  
824 tions: Pareto-optimal algorithm and its applications in energy markets,  
825 in: *e-Energy*, ACM, 2024, pp. 50–71.
- 826 [34] S. Angelopoulos, Online search with a hint, *Inf. Comput.* 295 (2023)  
827 105091.
- 828 [35] R. L. Rivest, A. R. Meyer, D. J. Kleitman, K. Winklmann, J. Spencer,  
829 Coping with errors in binary search procedures, *J. Comput. Syst. Sci.*  
830 20 (1980) 396–404.
- 831 [36] S. Im, R. Kumar, M. M. Qaem, M. Purohit, Online knapsack with

- 832 frequency predictions, in: Proc. Conference on Neural Information Pro-  
833 cessing Systems (NeurIPS), 2021, pp. 2733–2743.
- 834 [37] A. Kupavskii, E. Welzl, Lower bounds for searching robots, some faulty,  
835 Distributed Computing 34 (2019) 229–237.
- 836 [38] Y. Disser, S. Kratsch, Robust and adaptive search, in: Proc. Symposium  
837 on Theoretical Aspects of Computer Science (STACS), volume 66, 2017,  
838 pp. 26:1–26:14.
- 839 [39] S. Gal, Search Games, Academic Press, 1980.
- 840 [40] S. Angelopoulos, S. Kamali, D. Zhang, Online search with best-price  
841 and query-based predictions, in: Proc. AAAI Conference on Artificial  
842 Intelligence, 2022, pp. 9652–9660.
- 843 [41] S. Angelopoulos, S. Kamali, Rényi-Ulam games and online computation  
844 with imperfect advice, in: Proc. International Symposium on Math-  
845 ematical Foundations of Computer Science, (MFCS), 2023, pp. 13:1–  
846 13:15.
- 847 [42] A. López-Ortiz, S. Schuierer, On-line parallel heuristics, processor  
848 scheduling and robot searching under the competitive framework, The-  
849oretical Computer Science 310 (2004) 527–537.
- 850 [43] S. Angelopoulos, C. Dürr, S. Jin, Best-of-two-worlds analysis of on-  
851 line search, in: Proceedings of the 36th International Symposium on  
852 Theoretical Aspects of Computer Science (STACS), 2019, pp. 7:1–7:17.

853 [44] S. Gal, A general search game, Israel Journal of Mathematics 12 (1972)  
854 32–45.

855 [45] S. Alpern, S. Gal, The theory of search games and rendezvous, Kluwer  
856 Academic Publishers, 2003.

## 857 Appendix

### 858 Appendix A. Statement of Gal's theorem

859 Below, we give the formal statement of Gal's theorem.

860 **Theorem 26** (Gal [44]). *Let  $q$  be a positive integer, and  $X = (x_i)_{i=0}^{\infty}$  a*  
861 *sequence of positive numbers with  $\sup_{n \geq 0} x_{n+1}/x_n < \infty$  and  $\alpha_X > 0$ . Suppose*  
862 *that  $F_i$  is a sequence of functionals that satisfy the following properties:*

863 (1)  $F_i(X)$  depends only on  $x_0, x_1, \dots, x_{i+q}$ ,

864 (2)  $F_i(X)$  is continuous in every variable, for all positive sequences  $X$ ,

865 (3)  $F_i(aX) = F_i(X)$ , for all  $a > 0$ ,

866 (4)  $F_i(X + Y) \leq \max(F_i(X), F_i(Y))$ , for all positive sequences  $X, Y$ , and

867 (5)  $F_{i+j}(X) \geq F_i(X^{+j})$ , for all  $j \geq 1$ , where  $X^{+j} = (x_j, x_{j+1}, \dots)$ .

868 Then

$$\sup_{0 \leq q < \infty} F_q(X) \geq \sup_{0 \leq q < \infty} F_q(G_{\alpha_X}),$$

869 where  $G_a$  is defined as the geometric sequence  $(a^i)_{i=0}^{\infty}$ .

870 **Appendix B. Omitted proofs**

871 *Proof of Lemma 9.* For each  $j = 1, \dots, N$ , we define the function  $f_j(t) \in \mathbb{N}$ ,  
 872 where  $t > 0$ , as the number of  $x_{j,i}$ 's such that  $x_{j,i} \leq t$ . More precisely, for  
 873 any  $t > 0$ , the value  $f_j(t)$  is the unique nonnegative integer such that

$$x_{j,i} \leq t, \text{ for all } 1 \leq i \leq f_j(t) \quad \text{and} \quad x_{j,i} > t, \text{ for all } i > f_j(t).$$

874 In particular, if  $c_B i B^i \leq t$ , then  $f_j(t) \geq i$ . Therefore, if  $c_B i B^i \leq t <$   
 875  $c_B (i+1) B^{i+1}$ , using that  $\log(i+1) \leq \sqrt{i}$  for all  $i \geq 0$ , we conclude that

$$\begin{aligned} f_j(t) &\geq i > \frac{\log t - \log c_B}{\log B} - 1 - \frac{1}{\log B} \sqrt{i} \\ &\geq \frac{\log t - \log c_B}{\log B} - 1 - \frac{1}{\log B} f_j(t)^{\frac{1}{2}}. \end{aligned}$$

876 Similarly, noticing that, if  $t < c_A A^i$ , then  $f_j(t) < i$ , we conclude that

$$f_j(t) < i \leq \frac{\log t - \log c_A}{\log A} + 1,$$

877 whenever  $c_A A^{i-1} \leq t < c_A A^i$ .

878 Now, consider any large  $t > 0$  such that  $y_i \leq t < y_{i+1}$ , for some large  $i$ .

879 It must then hold that

$$i = \sum_{j=1}^N f_j(t),$$

880 whence

$$N \left( \frac{\log t - \log c_B}{\log B} - 1 - \frac{1}{\log B} \sqrt{i} \right) < i < N \left( \frac{\log t - \log c_A}{\log A} + 1 \right).$$

881

882 In particular, setting  $t = y_i$  yields that

$$c_A A^{\frac{i}{N}-1} < y_i < c_B e^{\sqrt{i}} B^{\frac{i}{N}+1},$$

883 which implies

$$\begin{aligned}
A^{\frac{1}{N}} &= \lim_{i \rightarrow \infty} c_A^{\frac{1}{i}} A^{\frac{1}{N} - \frac{1}{i}} \leq \liminf_{i \rightarrow \infty} y_i^{\frac{1}{i}} \leq \limsup_{i \rightarrow \infty} y_i^{\frac{1}{i}} \\
&\leq \lim_{i \rightarrow \infty} c_B^{\frac{1}{i}} e^{\frac{1}{\sqrt{i}}} B^{\frac{1}{N} + \frac{1}{i}} = B^{\frac{1}{N}},
\end{aligned}$$

884 thereby completing the proof of the lemma.  $\square$

885 *Proof of Lemma 23.* Consider a schedule  $X$  for RFT. For  $j \in [0, \dots, p-1]$ ,  
886 define  $l_X(t, j)$  as the length of the largest contract in  $S$  that has completed in  
887 processor  $j$  by time  $t$ . We also define by  $\ell_{X,f}(t)$  as the  $(f+1)$ -largest length  
888 in the set  $\{l_X(t, j)\}_{j=0}^{p-1}$ .

889 Following the notation of [6], we denote each contract  $c_j$  in  $X$  as a pair of  
890 the form  $(T_j, D_j)$ , where  $T_j$  is the start time of  $c_j$ , and  $D_j$  its length (as we  
891 will see, the specific processor to which the contract is assigned will not be  
892 significant for our analysis). We also define  $d_j$  to be equal to  $\ell_{X,f}(T_j + D_j - \epsilon)$ ,  
893 for  $\epsilon \rightarrow 0$ . In words,  $d_j$  is the longest contract length that has been completed  
894 right before  $c_j$  is about to terminate, assuming a worst-case scenario in which  
895  $f$  processors have been faulty, and they also happened to be the processors  
896 that have completed the longest contracts in the schedule by the said time:  
897 we call this contract length the  $(f+1)$  length relative to  $D_j$ .

898 Recall that  $\bar{X}$  denotes the sequence of all contract lengths in  $X$ , in non-  
899 decreasing order. Hence, each contract in  $X$  is mapped via its length to an  
900 element of this sequence (breaking ties arbitrarily).

901 Fix a time, say  $t$ , at which a contract  $c_{j_0} = (T_{j_0}, D_{j_0})$  terminates, say on  
902 processor 0 i.e.,  $t = T_{j_0} + D_{j_0}$ . For all  $m \in [1, p-1]$ , let  $c_{j_m} = (T_{j_m}, D_{j_m})$   
903 denote the longest contract length that has completed on processor  $m$  by  
904 time  $t$ . For every  $m \in [0, p-1]$ , define  $I_m$  as the set of indices in  $\mathbb{N}$  such

905 that  $i \in I_m$  if and only if the contract of length  $x_i$  has completed by time  $t$   
 906 in processor  $m$ . From the definition of the acceleration ratio, we have that

$$\text{acc}(X) \geq \frac{\sum_{i \in I_m} x_i}{d_{j_m}}, \quad \text{for all } m \in [0, p-1].$$

907 Therefore,

$$\text{acc}(X) \geq \max_{0 \leq m \leq p-1} \frac{\sum_{i \in I_m} x_i}{d_{j_m}},$$

908 and using the property  $\max\{a/b, c/d\} \geq \frac{a+b}{c+d}$ , for all  $a, b, c, d > 0$ , we obtain  
 909 that

$$\text{acc}(X) \geq \frac{\sum_{m=0}^{p-1} \sum_{i \in I_m} x_i}{\sum_{m=0}^{p-1} d_{j_m}}. \quad (\text{B.1})$$

910 Next, we will bound the numerator of the fraction in (B.1) from below and  
 911 its denominator from above. We begin with a useful observation: we can  
 912 assume, without loss of generality, that by time  $t$  (defined earlier), every  
 913 contract of length  $d_{j_0}$  or smaller has completed its execution. This follows  
 914 from the definition of  $d_{j_0}$ : if  $X$  completed a contract of length at most  $d_{j_0}$   
 915 later than time  $t$ , then one could simply “remove” this contract from  $X$ ,  
 916 and obtain a schedule of no worse acceleration ratio (in other words, such a  
 917 contract is useless, and one can derive a schedule of no larger acceleration  
 918 ratio than  $X$  that does not contain it).

919 Using the above observation, it follows that the numerator in (B.1) in-  
 920 cludes, as summands, all contracts of length at most  $d_{j_0}$ , as well as at least  
 921  $f + 1$  contracts that are at least as large as  $d_{j_0}$ . Let  $q$  denote an index such  
 922 that  $d_{j_0} = \bar{x}_q$ , then we have that

$$\sum_{m=0}^{p-1} \sum_{i \in I_m} x_i \geq \sum_{i=0}^{q+f+1} \bar{x}_i.$$



923 We now show how to upper-bound the denominator using the monotonicity  
 924 implied in the definition of the  $(f + 1)$  length relative to a given contract  
 925 length. Since  $c_{j_0}$  is completed no earlier than any other contract  $c_{j_m}$ , with  
 926  $m \in [1, p - 1]$ , we have that  $d_{j_m} \leq d_{j_0}$ . It thus follows that

$$\sum_{m=0}^{p-1} d_{j_m} \leq \sum_{i=q}^{q-(p-1)} \bar{x}_i.$$

927 Combining the two bounds, it follows that

$$\text{acc}(S) \geq \sup_{0 \leq q < \infty} \frac{\sum_{i=0}^{q+f+1} \bar{x}_i}{\sum_{i=q}^{q-(p-1)} \bar{x}_i}.$$

928 Define now the functional  $F_q(\bar{X}) = \frac{\sum_{i=0}^{q+f+1} \bar{x}_i}{\sum_{i=q}^{q-(p-1)} \bar{x}_i}$ , for every  $q$ . The functional  
 929 satisfies the conditions (1)-(5) of Theorem 26 (see Example 7.3 in [45]). More-  
 930 over,  $\sup_{n \geq 0} \bar{x}_{n+1}/\bar{x}_n < \infty$ , otherwise an infinite contract would be scheduled  
 931 in some processor, which, in turn, would render the corresponding processor  
 932 “useless”, since this contract would never complete. Last, we note that the  
 933 condition  $\alpha_X > 0$  is indeed satisfied, from Theorem 5 and Corollary 10.

934 By applying Gal’s Theorem (Theorem 26) it follows that

$$\text{acc}(X) \geq \sup_{0 \leq q < \infty} \frac{\sum_{i=0}^{q+f+1} \alpha_{\bar{X}}^i}{\sum_{i=q}^{q-(p-1)} \alpha_{\bar{X}}^i}.$$

935 If  $\alpha_{\bar{X}} \leq 1$ , then it is easy to show that the above expression shows that  
 936  $\text{acc}(X) = \infty$ ; see, e.g. [6]. Otherwise, i.e., if  $\alpha_{\bar{X}} > 1$ , after some simple  
 937 calculations along the lines of [6], we arrive at the desired result.  $\square$

### 938 **Appendix C. Additional experimental results**

939 In this section, we provide additional experimental results concerning the  
 940 RQS schedule.

941 *Appendix C.1. Experiments on the robustness  $r$*

942 In the main paper, we reported experiments for the setting in which  
943  $r = 4.0$ . Here, we compare RQS against the Oblivious and Baseline-Robust  
944 schedules for values of  $r$  such that  $r > 4$ . As before, the number of queries is  
945  $k = 16$ , and we assume  $\tau = 4 = k/4$  (i.e., the largest possible value of error  
946 the schedule can tolerate). The query error is generated as discussed in the  
947 main paper.

948 Figure C.5 depicts the experimental performance ratio for  $r \in \{4, 6, 8, 10\}$ .  
949 The results show that RQS is consistently better than Oblivious schedules  
950 and Baseline-Robust, in at least 95% of the possible interruption times.  
951 We also measured the average percentage improvement of RQS relative to  
952 Baseline-Robust: this is defined as the average, over all possible interrup-  
953 tions, of the signed % improvement in each interruption. These improve-  
954 ments are 11.88, 37.54, 42.96 and 45.48, for  $r \in \{4, 6, 8, 10\}$ , respectively.  
955 The experimental acceleration ratios of RQS are 2.39913, 1.88021, 1.83836,  
956 and 1.75828, respectively. We observe that the performance gains of RQS  
957 relative to Baseline-Robust increase as a function of  $r$ , as explained in the  
958 main paper.

959 In Figure C.5, we also plotted the minimum and maximum performance  
960 ratios of RQS, among the 100 random query responses, for each interruption  
961 time. As expected, the maximum performance ratio is similar to that of  
962 the Oblivious 1 schedule with a large base  $\zeta_{1,r}$ , since  $\tau = k/4$ ; however, this  
963 happens only very rarely, namely when the robust search algorithm finishes  
964 in a node that is very high up in the search tree. On average, we observe  
965 that RQS performs very close to the best run.

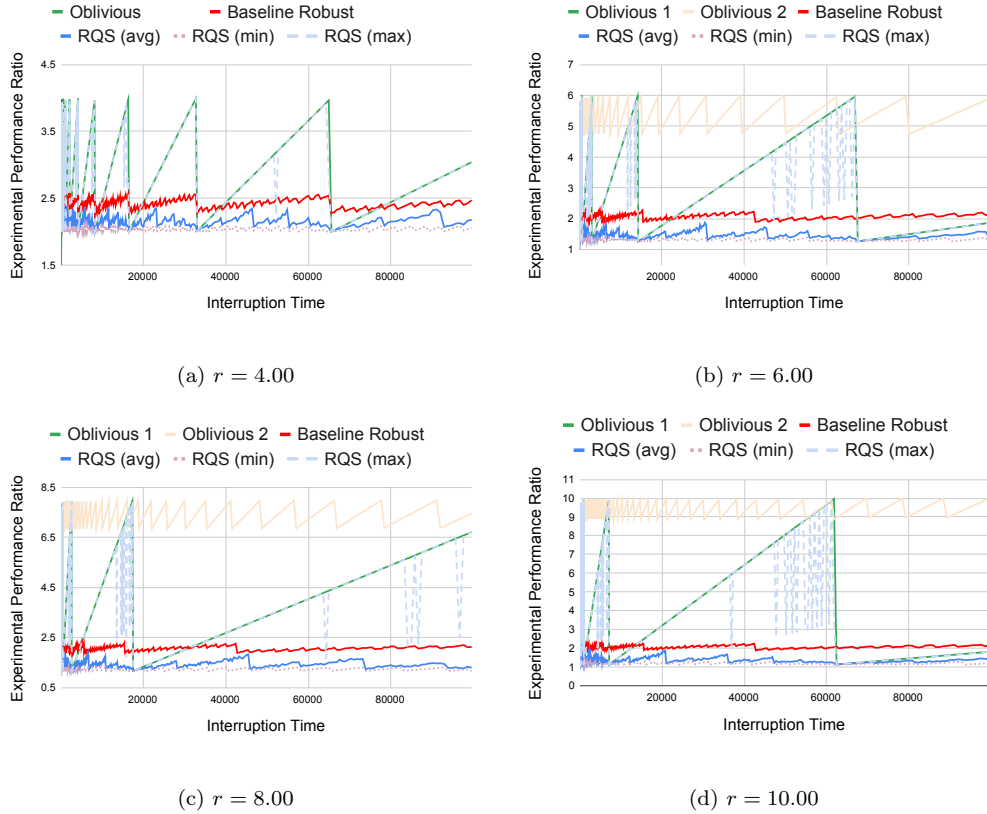


Figure C.5: Experimental performance ratios of various schedules for  $r \in \{4.00, 6.00, 8.00, 10.00\}$  when  $k = 16$  and  $\tau = 4$ .

966 *Appendix C.2. Experiments on the query error  $\eta$*

967 In the main paper, we considered the setting where the error takes values  
 968 from  $[0, \tau]$ . Here, we study a setting in which the maximum query error  
 969 can be either smaller or larger than the tolerance  $\tau$ . As before, we consider  
 970  $r = 4.00$ ,  $k = 16$ , and  $\tau = 4$ . Figure C.6 depicts the performance of schedules  
 971 when  $\eta$  query responses are randomly flipped, and  $\eta$  is uniformly distributed  
 972 in  $[0, \tau/2]$ ,  $[0, \tau]$ ,  $[0, 3\tau/2]$  and  $[0, 2\tau]$ . The last two ranges correspond to very  
 973 noisy query responses, for which the theoretical performance guarantees of

974 Theorem 18 do not necessarily hold.

975 We observe that, as expected, for relatively small values of  $\tau$ , and thus  
 976 small values of query error, RQS performs much better than Baseline-Robust.  
 977 This is because the robust search algorithm performs very well, if the error  
 978 is relatively small, and thus the schedule of RQS is very close to the Pareto-  
 979 optimal schedule of Theorem 18. If the error is too large, namely even if  $\eta$   
 980 exceeds  $\tau$ , RQS is still experimentally comparable to Baseline-Robust. Re-  
 981 call also our observation at the end of Section 4 concerning the worst-case  
 982 comparison of the two schedules.

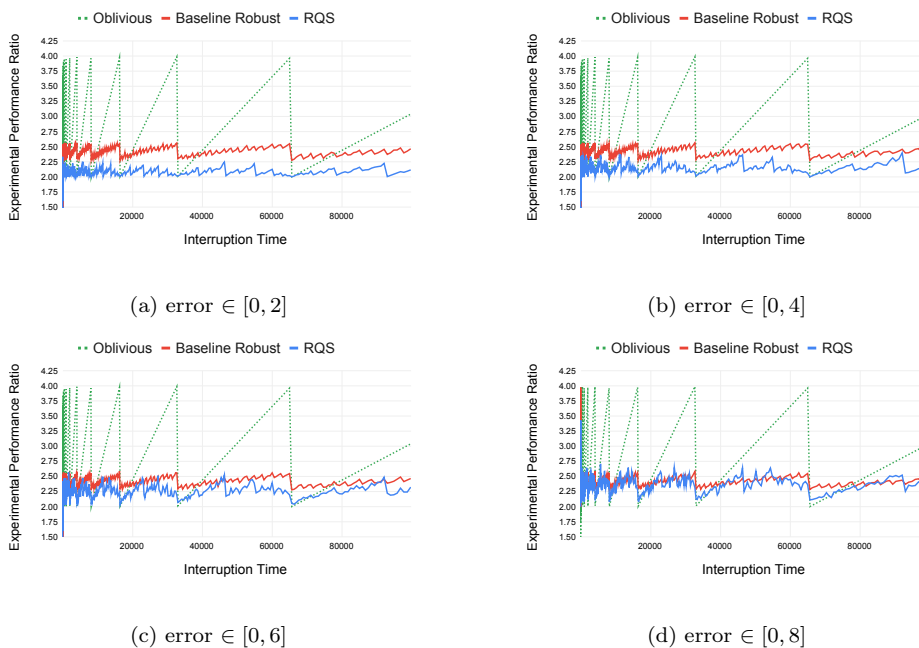




Figure C.6: Experimental performance ratio of various schedules for different ranges of error when  $r = 4.00$ ,  $k = 16$ , and  $\tau = 4$ .

983 *Appendix C.3. Experiments on the number of queries  $k$*


984 We report experiments for different values of the number of queries  $k$ , for  
985 the setting in which  $r = 4.00$ ,  $\tau = \lfloor k/4 \rfloor$  (note that we consider the high-  
986 est possible tolerance to errors). Figure C.7 summarizes our findings. We  
987 observe that RQS is consistently better than Baseline-Robust, and indepen-  
988 dently on the number of queries, in at least 95% of the possible interruption  
989 times. The average percentage improvements of RQS relative to Baseline-  
990 Robust are 8.75, 8.75, 11.50 and 11.51, for  $k \in \{8, 10, 12, 14\}$ , respectively.  
991 The experimental acceleration ratios of RQS are 2.045, 2.043, 2.167, and  
992 2.164 for  $k \in \{8, 10, 12, 14\}$ , respectively. We observe that RQS has a very  
993 stable performance, which is very close to the theoretically ideal acceleration  
994 ratio of 2, even for small constant  $k$ , in accordance with Theorem 20, and  
995 the discussion following the theorem in Section 4. For  $k = 12$ , both RQS  
996 and Baseline-Robust are somewhat more noisy: this is because  $k = 12$  is the  
997 value that maximizes the ratio  $k/\tau$  in this set of experiments.




(a)  $k = 8.00$



(b)  $k = 10.00$



(c)  $k = 12.00$



(d)  $k = 14.00$

Figure C.7: Experimental performance ratio of various schedules for  $k \in \{8, 10, 12, 14\}$  when  $r = 4.00$  and  $\tau = \lfloor k/4 \rfloor$ .