



Time-critical testing and search problems

This is a pre print version of the following article:

Original:

Agnētis, A., Hermans, B., Leus, R., Rostami, S. (2022). Time-critical testing and search problems. EUROPEAN JOURNAL OF OPERATIONAL RESEARCH, 296(2), 440-452 [10.1016/j.ejor.2021.03.038].

Availability:

This version is available <http://hdl.handle.net/11365/1191825> since 2022-03-02T20:21:04Z

Published:

DOI: <http://doi.org/10.1016/j.ejor.2021.03.038>

Terms of use:

Open Access

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. Works made available under a Creative Commons license can be used according to the terms and conditions of said license.

For all terms of use and more information see the publisher's website.

(Article begins on next page)

Time-critical testing and search problems

Alessandro Agnetis^a, Ben Hermans^{*,b}, Roel Leus^b, and Salim Rostami^c

^a*Dipartimento di Ingegneria dell'Informazione e Scienze Matematiche, Università di Siena, Italy*

^b*Research Center for Operations Research & Statistics, KU Leuven, Belgium*

^c*IÉSEG School of Management, France*

Abstract

This paper introduces a problem in which the state of a system needs to be determined through costly tests of its components by a limited number of testing units and before a given deadline. We also consider a closely related search problem in which there are multiple searchers to find a target before a given deadline. These natural generalizations of the classical sequential testing problem and search problem are applicable in a wide range of time-critical operations such as machine maintenance, diagnosing a patient, and new product development. We show that both problems are NP-hard, develop a pseudo-polynomial dynamic program for the special case of two time slots, and describe a partial-order-based as well as an assignment-based mixed integer program for the general case. Based on extensive computational experiments, we find that the assignment-based formulation performs better than the partial-order-based formulation for the testing variant, but that this is the other way round for the search variant. Finally, we propose a pairwise-interchange-based local search procedure and show that, empirically, it performs very well in finding near-optimal solutions.

1 Introduction

The timely diagnosis of a system constitutes a key task in a wide range of time-critical operations. The downtime of a machine due to a periodic or corrective maintenance, for example, crucially depends on the time needed to test which components are functioning correctly. Medical tests oftentimes also have a high time criticality, not only to establish a correct diagnosis, but also to identify an effective treatment. When developing a new product, in turn, it is essential to complete all tests concerning the product's safety and performance in time, because a delayed product launch could enable competitors to grasp a leading market share with a similar product.

Each of the above settings can be generically modeled as a multi-component system whose state is either up or down (healthy or unhealthy; safe or unsafe) depending on which of its n components are functioning (Ünlüyurt, 2004). Consider, for example, a machine that can only operate if all its components are functioning, a patient who only passes a medical examination if all test results are negative, or a new product that can only be launched if all safety and performance tests prove to be successful. These are all examples of so-called *serial* or *n -out-of- n systems*, in which the state is only up if all of its components are functioning. A *parallel* or *1-out-of- n system*, on the other hand, is up if there is at least one functioning component. For instance, it suffices to find a single treatment that works to cure a patient.

*The author is funded by a Postdoctoral Fellowship of the Research Foundation – Flanders.

The *sequential testing problem* deals with checking the state of a system through costly tests of its components, where each component has a given failure probability (Ünlüyurt, 2004). Failures are typically assumed to occur independently with a probability that is known from, for example, the analysis of historical data. The cost to test a component can measure both monetary expenses as well as intangible aspects such as how inconvenient the test is to a patient. Once the system’s state is known, the testing procedure halts. The goal is then to determine in which sequence to test the different components so as to minimize the total expected testing cost.

In this paper, we consider a variant of the classical sequential testing problem in which every test takes unit time and in which the system’s state has to be known within a given deadline T . In order to comply with such timeliness requirements, there are m testing units on which tests can take place simultaneously, and the testing procedure is thus no longer purely sequential. As soon as the system’s state is known, all testing units are halted. This leads to what we call the *time-critical testing problem*, where the goal is to schedule the n tests on the m available testing units so as to determine the system’s state within the deadline T at minimum expected cost. Although we focus on serial systems, all our results are also directly applicable to parallel systems. Indeed, if we invert the interpretation of a test’s outcome and of the definition of a system’s up and down state, then the testing of a 1-out-of- n system becomes equivalent to the testing of a n -out-of- n system (Ünlüyurt, 2004).

We also study a closely related *time-critical search problem* in which there are m searchers and n possible locations that could contain a given target, and the objective is to find the target before a deadline T at minimum expected cost. This is a variant of the time-critical testing problem where the tests are dependent in the sense that exactly one of the system’s components is defective, and where we need to identify this component (Wagner and Davis, 2001). Applications to this problem include a malfunctioning component of a machine that needs to be identified before a certain deadline, a terrorist that needs to be captured before an attack takes place, or a person in distress who needs to be found in the context of a search and rescue operation.

The main contributions of this paper are as follows. We formally define the time-critical testing problem and the time-critical search problem, and establish that both problems are strongly NP-hard in general. We show in particular that the time-critical search problem reduces to the time-critical testing problem, which extends an earlier result of Kelly (1982) towards multiple testing units. For the special case where $T = 2$, we establish weak NP-hardness and describe a pseudo-polynomial dynamic program as well as a fully polynomial-time approximation scheme (FPTAS). Next, we introduce a partial-order-based and an assignment-based mixed integer programming (MIP) formulation for each of the problem variants, and show how to adapt these formulations to solve the precedence-constrained sequential testing problem (Monma and Sidney, 1979) and the batch-testing problem (Daldal et al., 2016). To the best of our knowledge, no compact exact MIP formulations were previously available for the latter two problems. Finally, we illustrate that intuitive approaches such as a greedy heuristic or a pairwise-interchange-based local search procedure fail to provide a constant-factor approximation guarantee.

Based on extensive computational experiments, we find that the assignment-based formulation performs better than the partial-order-based formulation for the testing variant, while this is the other way round for the search variant. Moreover, instances for the time-critical testing problem are considerably harder to solve than those for the search variant. Despite its lack of a theoretical worst-case guarantee, we also find that, empirically, our local search procedure seems to perform very well in finding near-optimal solutions in limited computation times.

The remainder of this paper is organized as follows. After the literature review in Section 2, we formally define the time-critical testing problem and search problem in Section 3. Next, we examine the structural properties of an optimal schedule in Section 4, analyze the problems’

complexity in Section 5, and consider the special case of two time slots in Section 6. Sections 7 and 8, in turn, discuss our proposed MIP formulations and local search heuristic, whose performance is evaluated in the computational experiments reported in Section 9. Section 10, finally, concludes and indicates directions for further research.

2 Related work

Since the U.S. Air Force highlighted the importance of the sequential testing problem (Johnson, 1956), it has received considerable attention in the literature. For a serial system without precedence constraints, the following intuitive result has been established independently in a variety of different contexts (Mitten, 1960; Boothroyd, 1960; Kadane, 1969): every sequence that inspects the components in non-decreasing order of their cost-to-failure-probability ratio is optimal. This idea has subsequently been generalized into polynomial-time algorithms for the sequential testing problem with series-parallel precedence constraints (Garey, 1973; Monma and Sidney, 1979; Berend et al., 2014). For general precedence constraints, the problem is NP-hard (Kelly, 1982; De Reyck and Leus, 2008); Rostami et al. (2019) describe a branch-and-price algorithm as well as a dynamic program for this case. We refer to Ünliyurt (2004) for an excellent and detailed review of the sequential testing literature.

Most closely related to our work is the *batch-testing problem* for serial systems as introduced by Daldal et al. (2016). In this variant of the sequential testing problem, tests can be performed simultaneously in batches and there is a fixed set-up cost per batch. Daldal et al. (2016) develop an approximation algorithm for the problem with a $6.829 + \varepsilon$ worst-case guarantee factor and an integer program that approximates the optimal solution within a factor $1 + \varepsilon$, where $\varepsilon \in (0, 1)$ is arbitrary but fixed. Daldal et al. (2017) study the same problem with additional restrictions on the set of tests that can be combined in a batch, and Segev and Shaposhnik (2018) use techniques from approximate dynamic programming to devise a polynomial-time approximation scheme (PTAS) for the batch-testing problem. Contrary to the aforementioned articles, we consider a deadline, an upper bound m on the batch size, and no fixed set-up cost per batch. For this reason, the approximation algorithms as proposed by Daldal et al. (2016) and Segev and Shaposhnik (2018) do not seem to be directly modifiable towards our setting. The assignment-based formulation to be described in Section 7, however, can be adapted such that it also solves the batch-testing problem.

Agnetis et al. (2009) define a problem where n unreliable jobs with given revenues are to be processed on m machines. If a job fails, then all subsequent jobs on the same machine are blocked, and their revenue is lost. The goal is then to assign the jobs to machines so as to maximize the expected revenue. Agnetis et al. (2009) give a polyhedral characterization for the case of one machine and show that the problem is NP-hard for two machines. In subsequent articles, Agnetis et al. (2014) show that a list scheduling algorithm based on the so-called Z-ratio leads to a 0.8535-approximation for the case of two machines, and to a 0.8531-approximation in the general case of m machines (Agnetis and Lidbetter, 2020). Our work differs from the aforementioned articles since we minimize the expected cost rather than maximize the expected revenue, and since we assume that the testing procedure halts on all units as soon as the system's state is known.

Another related problem is the *salvo policy problem*, in which there are a limited number of moments in time (the so-called salvos) at which multiple missiles can be fired simultaneously to stop an incoming object (Gould, 1984). If all the missiles fired at a certain salvo are unsuccessful, the missiles scheduled at the next salvo must be fired. Each missile fired at a certain salvo has a hit probability that only depends on the salvo, and the goal is to minimize the expected number

of missiles that need to be fired. Van Ee (2020) develops algorithms to solve different variants of the problem, and we refer to his article as well as the work of Glazebrook and Washburn (2004) for a more detailed literature review. The main difference with our work is that, in our setting, the failure probabilities and testing costs are time-independent and can differ between components. Hence, we can also interpret the time-critical testing problem as a variant of the salvo policy problem where there are T salvos and n missiles with different time-independent costs and hit probabilities, and where no more than m missiles can be fired at each salvo. To the best of our knowledge, this problem variant has not been studied before.

If we identify success probabilities with job weights and search costs with job processing times, then the classical sequential search problem (i.e., the time-critical search problem with a single searcher and no deadline) is equivalent to the problem $1 \parallel \sum w_j C_j$ of minimizing the total weighted completion time of a job set on a single machine. Interestingly, numerous results for $1 \parallel \sum w_j C_j$ parallel those for the sequential testing problem: without precedence constraints it is optimal to search the different locations in non-decreasing order of their cost-to-probability ratio (Smith, 1956; Gluss, 1959), this rule can be generalized to solve the case with series-parallel precedence constraints efficiently (Lawler, 1978; Sidney, 1975), and the problem is strongly NP-hard for general precedence constraints (Lenstra and Rinnooy Kan, 1978). In fact, Kelly (1982) shows that the sequential testing problem is at least as hard as the problem $1 \parallel \sum w_j C_j$. Our results generalize these findings to a time-critical setting with multiple testing units, since we show in Section 5 that the time-critical search problem reduces to the time-critical testing problem.

In the rich literature on search theory (see for instance the books of Stone (2007) or Alpern et al. (2013) for a detailed overview), there are a number of articles that consider search problems with multiple searchers (Li and Huang, 2018), general cost and weight functions (Fokkink et al., 2019; Happach et al., 2020), or an unknown deadline (Lin and Singham, 2016; Lidbetter, 2020). These problems, however, are overall quite different from our setting, and the time-critical search problem as defined in this paper appears to be novel.

3 Problem definition and notation

Consider an n -out-of- n system whose state is up if and only if all of its n components are functioning; otherwise, the system is down. Each component $j \in N := \{1, \dots, n\}$ is functioning with a given *success probability* $p_j \in [0, 1] \cap \mathbb{Q}$, where this event is independent of the outcome for other components, and we can determine whether a component is functioning by means of a test with *testing cost* $c_j \in \mathbb{N}$. We assume that a test requires one time unit, that the cost of a test is borne regardless of its outcome, and that tests are perfect in the sense that neither type-1 nor type-2 errors may occur. There are $m \leq n$ identical *testing units* or *machines* on which tests can take place simultaneously, and we require that all tests be completed within a *deadline* $T \leq n$. To ensure feasibility, we assume that $n \leq mT$. If a test fails, then the process stops on all machines and we conclude that the system is down. The objective is to schedule all n tests on the m available machines so as to determine the system's state at minimum expected cost within the deadline T .

Since all tests require unit time, we can divide the time horizon into T equal *time slots*. Moreover, since the testing procedure is halted on all machines as soon as the system's state is known, it is not relevant which machine a given test is assigned to. Hence, we can represent a feasible solution or *schedule* by an ordered partition $\sigma = (S_1, \dots, S_T)$ of the set N into T sets, where $S_t \subseteq N$ with $|S_t| \leq m$ denotes the set of tests scheduled in time slot $t = 1, \dots, T$. For an arbitrary schedule $\sigma = (S_1, \dots, S_T)$, denote by $\mathcal{A}_t := \bigcup_{k=1}^{t-1} S_k$ the set of all tests scheduled before

time slot $t \in \{1, \dots, T\}$, where $\mathcal{A}_1 = \emptyset$. Given a subset $S \subseteq N$ and a sequence $(a_j)_{j \in N}$ of real numbers, we use the conventional notation that $a(S) := \sum_{j \in S} a_j$, and we let $\sum_{j \in \emptyset} a_j = 0$ and $\prod_{j \in \emptyset} a_j = 1$. The *time-critical testing problem* (TCTP) then asks for a schedule $\sigma = (S_1, \dots, S_T)$ that minimizes the *expected testing cost*

$$z(\sigma) := \sum_{t=1}^T c(S_t) \prod_{j \in \mathcal{A}_t} p_j. \quad (1)$$

Here, each index t associates the cost $c(S_t)$ with the probability that the tests in S_t need to be performed because all components tested before time t were successful (and the system's state is thus still unknown). A TCTP instance is completely defined by the tuple $(m, n, T, (c_j, p_j)_{j \in N})$, and we refer to m -TCTP as the special case of the problem when the number of machines is fixed to m .

We also consider a search problem related to TCTP, where a given target is hidden in *exactly one* out of n possible locations. Each location $j \in N := \{1, \dots, n\}$ has a probability $\pi_j \in [0, 1] \cap \mathbb{Q}$ to contain the target, where $\sum_{j \in N} \pi_j = 1$. Searching location j takes unit time and leads to a cost $c_j \in \mathbb{N}$. Given that there are $m \leq n$ available searchers and that the object needs to be found before a certain deadline $T \leq n$, where $n \leq Tm$, the question is when to search each location so as to minimize the expected cost. Using the same notation as above, a feasible schedule can be represented by an ordered partition $\sigma = (S_1, \dots, S_T)$ of the set N into T sets with $|S_t| \leq m$ for each $t = 1, \dots, T$. The *time-critical search problem* (TCSP) then asks for a schedule $\sigma = (S_1, \dots, S_T)$ that minimizes the *expected search cost*

$$\bar{z}(\sigma) := \sum_{t=1}^T c(S_t) \sum_{k=t}^T \pi(S_k). \quad (2)$$

Here, the index t in the outer summation reflects the fact that we incur cost $c(S_t)$ when the target is in one of the locations searched in time slots $k = t, \dots, T$, which occurs with probability $\sum_{k=t}^T \pi(S_k)$. A TCSP instance is completely defined by the tuple $(m, n, T, (c_j, \pi_j)_{j \in N})$, and we refer to m -TCSP as the special case with a fixed number of m searchers.

The difference between TCTP and TCSP consists in the way in which both variants model the probability that a job (i.e., testing a component or searching a location) needs to be performed. While for TCTP the success probabilities are taken to be independent, we assume for TCSP that exactly one location contains the target and the probabilities thus add up to one. For TCTP, the probability that a job needs to be performed then equals the product of success probabilities over all jobs scheduled in preceding time slots. For TCSP, in turn, this probability equals one minus the probability that one of the locations searched in preceding time slots contains the target.

4 Structural properties

If there is only a single testing unit or searcher, then TCTP and TCSP reduce to the sequential testing problem and to the problem 1 || $\sum w_j C_j$ of minimizing the weighted completion time on a single machine, respectively. Hence, as mentioned in Section 2, every sequence that tests the components or searches the locations in non-decreasing order of, respectively, the ratio $c_j/(1-p_j)$ or c_j/π_j is optimal. Lemma 1, stated below, implies that the resulting schedule is also optimal for the setting with multiple machines or searchers if the deadline allows for this. That is, the only reason to perform tests or to search locations simultaneously is to comply with a deadline $T < n$.

Lemma 1. *For each TCTP or TCSP instance there exists an optimal schedule $\sigma^* = (S_1^*, \dots, S_T^*)$ in which S_t^* is non-empty for all $t \leq T$.*

Proof. Since the proof is completely analogous for TCTP and TCSP, we only consider the former case. For an arbitrary TCTP instance, let $\sigma = (S_1, \dots, S_T)$ be an optimal schedule that does not satisfy the property. By a straightforward pairwise interchange argument on Equation (1), we can assume without loss of generality that all non-empty sets are scheduled before the empty sets in σ . Call t the first time slot for which $|S_t| \geq 2$; such a slot must exist because otherwise σ would satisfy the property (recall that our definition of a TCTP instance assumes that $T \leq n$). Now take an arbitrary test $j \in S_t$ and define a new schedule $\sigma' = (S_1, \dots, S_t \setminus \{j\}, \{j\}, S_{t+1}, \dots, S_{T-1})$. Equation (1) then yields that

$$z(\sigma) - z(\sigma') = \left(\prod_{i \in \mathcal{A}_t} p_i \right) \left(1 - \prod_{i \in S_t \setminus \{j\}} p_i \right) c_j \geq 0.$$

Repeating this argument at most n times yields an optimal schedule that satisfies the property. \square

The following lemma generalizes the ratio rule to the setting of multiple machines. In particular, given a subset $S \subseteq N$ for a TCTP instance, define the *cost-to-probability ratio* as

$$\rho(S) := \frac{c(S)}{1 - \prod_{j \in S} p_j}$$

if $\prod_{j \in S} p_j \neq 1$, and $\rho(S) := +\infty$ otherwise. For an instance of TCSP, in turn, we define the ratio

$$\bar{\rho}(S) := \frac{c(S)}{\pi(S)}$$

if $\pi(S) \neq 0$, and $\bar{\rho}(S) := +\infty$ otherwise. Lemma 2 below implies that, for both TCTP and TCSP, there exists an optimal schedule that sequences the different sets in non-decreasing order of the corresponding ratios.

Lemma 2. *Reordering the sets S_t of a schedule $\sigma = (S_1, \dots, S_T)$ in non-decreasing order of the ratio $\rho(S_t)$ or $\bar{\rho}(S_t)$ does not increase the expected testing or search cost, respectively.*

Proof. Since the proof is again completely analogous for TCTP and TCSP, we only consider the former case. Let $\sigma = (S_1, \dots, S_T)$ be a schedule with $\rho(S_t) > \rho(S_{t+1})$ for some $t < T$ and define the schedule σ' from σ in which S_t and S_{t+1} have been interchanged. Equation (1) then yields that

$$z(\sigma) - z(\sigma') = \left(\prod_{j \in \mathcal{A}_t} p_j \right) \left[\left(1 - \prod_{j \in S_{t+1}} p_j \right) c(S_t) - \left(1 - \prod_{j \in S_t} p_j \right) c(S_{t+1}) \right],$$

which is non-negative. Repeating this procedure at most T^2 times then yields the result. \square

Based on Lemma 2, one might propose a greedy approach where we test in each time slot a subset S of the remaining components for which the ratio $\rho(S)$ is minimal among all sets that contain sufficient tests to meet the deadline. The next example, however, shows that this approach does not lead to a constant-factor approximation guarantee for TCTP. A similar example shows that an analogous greedy approach that minimizes the ratio $\bar{\rho}(S)$ also fails to provide a constant-factor approximation guarantee for TCSP.

Example 1. Consider a TCTP instance with $n = 3$ tests, $m = 2$ machines, and deadline $T = 2$. Given an integer $M > 0$, we define the costs and success probability of each job as follows: $c_1 = 1$, $p_1 = 1/M$, $c_2 = 0$, $p_2 = 1 - 1/M$, $c_3 = M$, and $p_3 = 1 - 1/M$. The greedy approach described above would then test component 2 in the first time slot (its ratio equal to 0 is clearly minimal), and the remaining components $\{1, 3\}$ in the second time slot. This leads to an expected testing cost equal to $z_{\text{greedy}} = c_2 + p_2(c_1 + c_3) = 0 + (1 - 1/M)(M + 1) = M - 1/M$. The schedule that tests $\{1\}$ first and $\{2, 3\}$ next, however, attains an expected testing cost equal to $z^* = c_1 + p_1(c_2 + c_3) = 1 + (1/M)M = 2$. Hence, the ratio z_{greedy}/z^* is $\Omega(M)$ and thus cannot be bounded by a constant.

The next lemma reveals that, for TCTP, the range between the minimal and maximal expected testing cost, and thus the potential benefit of identifying an optimal schedule, increases as the joint success probability $\prod_{j \in N} p_j$ decreases. Intuitively, if the system has a higher probability to be up, then it becomes more likely that all tests need to be performed, and their specific ordering becomes less influential. Observe that an analogous result does not hold for TCSP, since the probability $\sum_{j \in N} \pi_j$ is fixed to one for this problem.

Lemma 3. *Given a TCTP instance, let $\sigma^* = (S_1^*, \dots, S_T^*)$ be an optimal schedule. For every schedule $\sigma = (S_1, \dots, S_T)$ it then holds that $z(\sigma^*) \geq z(\sigma) \prod_{j \in N} p_j$.*

Proof. Consider an arbitrary schedule $\sigma = (S_1, \dots, S_T)$. Since $z(\sigma) \leq c(N)$, we obtain that

$$z(\sigma^*) = \sum_{t=1}^T c(S_t^*) \prod_{j \in \mathcal{A}_t^*} p_j \geq \sum_{t=1}^T c(S_t^*) \prod_{j \in N} p_j = c(N) \prod_{j \in N} p_j \geq z(\sigma) \prod_{j \in N} p_j. \quad \square$$

5 Complexity analysis

In this section we analyze the complexity of TCTP and TCSP. First, we show in Section 5.1 that TCSP is NP-hard even if $T = 2$ and that the m -TCSP is strongly NP-hard for every $m \geq 3$. Next, we show that TCSP reduces to TCTP in Section 5.2, so that the hardness results for TCSP carry over directly to TCTP.

5.1 Hardness of TCSP

Given a constant $\alpha \in \mathbb{Q}$, consider the decision version of TCSP, which we call TCSP-D, that asks whether there exists a schedule σ such that $\bar{z}(\sigma) \leq \alpha$. To establish the complexity of TCSP-D, we will use the following lemma:

Lemma 4. *For every $r \in \mathbb{N}$ numbers $a_1, \dots, a_r \in \mathbb{R}$ with a fixed sum $A := \sum_{i=1}^r a_i$ it holds that*

$$\sum_{i=1}^r \sum_{j=i}^r a_i a_j \geq \frac{(r+1)A^2}{2r},$$

with equality holding if and only if $a_i = A/r$ for every $i = 1, \dots, r$.

Proof. For arbitrary numbers $r \in \mathbb{N}$ and $A \in \mathbb{R}$, define the set $X := \{(x_1, \dots, x_r) \in \mathbb{R}^r : \sum_{i=1}^r x_i = A\}$ and the function $f: X \rightarrow \mathbb{R}: (x_1, \dots, x_r) \mapsto \sum_{i=1}^r \sum_{j=i}^r x_i x_j$. It then suffices to show that f attains its unique minimum if $x_i = A/r$ for each $i = 1, \dots, r$. To see that this is the case, observe that the function $g: \mathbb{R}^{r-1} \rightarrow \mathbb{R}: (x_2, \dots, x_r) \mapsto f(A - \sum_{i=2}^r x_i, x_2, \dots, x_r)$ is strictly convex with partial derivatives equal to zero if and only if $x_j = A - \sum_{i=2}^r x_i$ for each $j = 2, \dots, r$. \square

TCSP-D is in NP since, for a given schedule σ , we can use Equality (2) to verify in polynomial time whether $\bar{z}(\sigma) \leq \alpha$. The next theorem uses a reduction from the NP-complete PARTITION problem (Garey and Johnson, 1979) to show that TCSP-D is NP-complete even for the special case of two time slots.

Theorem 1. *TCSP-D is NP-complete, even if $T = 2$.*

Proof. Consider an instance of the PARTITION problem where, given $q \in \mathbb{N}$ natural numbers $u_1, \dots, u_q \in \mathbb{N}$ with $A := \sum_{i=1}^q u_i$, the question is whether there exists a set $S \subseteq \{1, \dots, q\}$ such that $u(S) = A/2$. Now define a TCSP-D instance with $m = q$ searchers, deadline $T = 2$, and $n = q$ locations with $c_j = u_j$ and $\pi_j = u_j/A$ for all $j = 1, \dots, q$. Finally, we let $\alpha = 3A/4$.

For every schedule $\sigma = (S_1, S_2)$, Lemma 4 (with $r = 2$ and identifying a_i with $u(S_i)$ for $i = 1, 2$) yields that

$$\sum_{t=1}^2 c(S_t) \sum_{k=t}^2 \pi(S_k) = \frac{1}{A} \sum_{t=1}^2 \sum_{k=t}^2 u(S_t)u(S_k) \geq \frac{1}{A} \frac{3A^2}{4} = \frac{3A}{4} = \alpha.$$

Hence, the answer to the TCSP-D instance is ‘yes’ if and only if there exists a schedule $\sigma = (S_1, S_2)$ for which the above inequality holds with equality. From Lemma 4, this occurs if and only if $u(S_1) = u(S_2) = A/2$, i.e., if and only if the answer to the PARTITION instance is ‘yes’. \square

The following theorem shows that TCSP-D, where T is part of the input, is strongly NP-complete for every fixed number $m \geq 3$ of searchers. The proof uses a reduction from the strongly NP-complete 3-PARTITION problem (Garey and Johnson, 1979).

Theorem 2. *TCSP-D with a fixed number $m \geq 3$ of searchers is strongly NP-complete.*

Proof. The proof uses a reduction from the 3-PARTITION problem. Given $q, B \in \mathbb{N}$, consider $3q$ integers u_1, \dots, u_{3q} with $\sum_{j=1}^{3q} u_j = qB$ and $B/4 < u_j < B/2$ for each $j = 1, \dots, 3q$. The 3-PARTITION problem then asks whether $\{u_1, \dots, u_{3q}\}$ can be partitioned into q subsets U_1, \dots, U_q with equal sum $u(U_t) = B$ for each $t = 1, \dots, q$. Observe that, since $B/4 < u_i < B/2$ for each $i = 1, \dots, 3q$, the subset U_t must have equal cardinality $|U_t| = 3$ for each $t = 1, \dots, q$.

For an arbitrary 3-PARTITION instance $(q, B, u_1, \dots, u_{3q})$, define a TCSP-D instance with $\alpha = (q+1)B/2$, a fixed number of $m \geq 3$ searchers, a deadline $T = q$, and $n = 3q$ locations with $c_j = u_j$ and $\pi_j = u_j/(qB)$ for every $j \in N$.

For every schedule $\sigma = (S_1, \dots, S_T)$ it follows from Lemma 4 (by taking $r = T = q$, and by identifying the numbers a_i with $u(S_i)$ for $i = 1, \dots, T$) that

$$\sum_{t=1}^T c(S_t) \sum_{k=t}^T \pi(S_k) = \frac{1}{qB} \sum_{t=1}^T \sum_{k=t}^T u(S_t)u(S_k) \geq \frac{1}{qB} \frac{(q+1)(qB)^2}{2q} = \frac{(q+1)B}{2} = \alpha.$$

Hence, the answer to the TCSP-D instance is ‘yes’ if and only if there exists a schedule $\sigma = (S_1, \dots, S_T)$ for which the above inequality holds at equality. From Lemma 4, this occurs if and only if $u(S_t) = qB/q = B$ for all $t = 1, \dots, T$. Observe that, since $B/4 < u_i < B/2$ for each $i = 1, \dots, 3q$, this latter condition can only be satisfied if σ schedules exactly three jobs in every time slot, such that the schedule is feasible with respect to the machine capacity $m \geq 3$. \square

5.2 Hardness of TCTP

We define the decision version TCTP-D of TCTP as follows: given a constant $\beta \in \mathbb{Q}$, does there exist a schedule σ such that $z(\sigma) < \beta$? To show that TCSP-D reduces to TCTP-D in Theorem 3,

we apply a similar approach as the one used in Hermans et al. (2019) to decompose and to bound a product of probabilities. In particular, we need the following lemma, which can be shown by induction.

Lemma 5 (Hermans et al., 2019). *Let $r \in \mathbb{N}$ and $a_1, \dots, a_r \in \mathbb{R}$, then*

$$\prod_{i=1}^r (1 - a_i) = 1 - \sum_{i=1}^r a_i + \sum_{i=1}^r \sum_{j=i+1}^r a_i a_j - \prod_{k=j+1}^r (1 - a_k).$$

Theorem 3. *TCSP-D reduces to TCTP-D.*

Proof. Given an instance of TCSP-D $(\alpha, m, n, T, (c_j, \pi_j)_{j \in N})$, let $W \in \mathbb{N}$ be the smallest integer such that $w_j := \pi_j W$ is integer for all $j \in N$. Since all probabilities $(\pi_j)_{j \in N}$ are rational numbers, W is polynomially bounded in the input size. Define an associated TCTP-D instance with $p_j = 1 - w_j/M$ for every $j \in N$, where $M := c(N)(nW)^2$, and with all other parameters equal to the corresponding ones of the TCSP-D instance. Moreover, let $\gamma := (c(N) - \alpha)W$ and $\beta = c(N) - (\gamma - 1)/M$. We want to show that there exists a schedule σ with $\bar{z}(\sigma) \leq \alpha$ if and only if there exists a schedule σ for the TCTP instance such that $z(\sigma) < \beta$.

Let $\sigma = (S_1, \dots, S_T)$ be an arbitrary schedule, then it follows from Equation (1) and the definition of $(p_j)_{j \in N}$ that

$$z(\sigma) = \sum_{t=1}^T c(S_t) \prod_{j \in \mathcal{A}_t} \left(1 - \frac{w_j}{M}\right). \quad (3)$$

Now define for every set $S \subseteq N$ the number

$$Q(S) = \sum_{i \in S} \sum_{j \in S: j > i} \frac{w_i w_j}{M^2} \prod_{k \in S: k > j} \left(1 - \frac{w_k}{M}\right), \quad (4)$$

and note that, by choice of M , it holds for each $i, j, k \in S$ that

$$0 \leq \frac{w_i w_j}{M} < \frac{1}{c(N)n^2} \quad \text{and} \quad 0 < 1 - \frac{w_k}{M} \leq 1.$$

Since Equation (4) contains strictly less than n^2 terms, this implies that $0 \leq Q(S) < 1/(Mc(N))$ for each $S \subseteq N$. Hence, by using Lemma 5 to rewrite Equation (3), we obtain that

$$\begin{aligned} z(\sigma) &= \sum_{t=1}^T c(S_t) \left(1 - \frac{1}{M} w(\mathcal{A}_t) + Q(\mathcal{A}_t)\right) \\ &< c(N) + \frac{1}{M} - \frac{1}{M} \sum_{t=2}^T c(S_t) w(\mathcal{A}_t) \end{aligned} \quad (5)$$

and

$$z(\sigma) \geq c(N) - \frac{1}{M} \sum_{t=2}^T c(S_t) w(\mathcal{A}_t). \quad (6)$$

If $\sum_{t=2}^T c(S_t) w(\mathcal{A}_t) \geq \gamma$, then Inequality (5) implies that $z(\sigma) < c(N) - (\gamma - 1)/M = \beta$. Conversely, if $\sum_{t=2}^T c(S_t) w(\mathcal{A}_t) < \gamma$, then we also know that $\sum_{t=2}^T c(S_t) w(\mathcal{A}_t) \leq \gamma - 1$ since $(c_j)_{j \in N}$ and $(w_j)_{j \in N}$ are integer. Inequality (6) then yields that $z(\sigma) \geq c(N) - (\gamma - 1)/M = \beta$. Hence, we obtain that

$$\sum_{t=2}^T c(S_t) \sum_{k=1}^{t-1} w(S_k) = \sum_{t=2}^T c(S_t) w(\mathcal{A}_t) \geq \gamma = (c(N) - \alpha)W \quad (7)$$

if and only if $z(\sigma) < c(N) - (\gamma - 1)/M = \beta$.

Finally, Equation (2) combined with $\pi(N) = \sum_{k=1}^T \pi(S_k) = 1$ and $w_j = \pi_j W$ yields that

$$(c(N) - \bar{z}(\sigma)) W = \sum_{t=1}^T c(S_t) \left(\sum_{k=1}^T \pi(S_k) - \sum_{k=t}^T \pi(S_k) \right) W = \sum_{t=2}^T c(S_t) \sum_{k=1}^{t-1} w(S_k).$$

A schedule σ thus satisfies Inequality (7) if and only if $\bar{z}(\sigma) \leq \alpha$, which completes the proof. \square

TCTP-D is in NP since, for a given schedule σ , we can use Equality (1) to verify in polynomial time whether $z(\sigma) < \beta$. Hence, together with Theorems 1, 2, and 3, we obtain this section's main result.

Corollary 1. *TCTP-D is NP-complete even if $T = 2$. If T is part of the input, then the problem is strongly NP-complete for every fixed number $m \geq 3$ of machines.*

6 Special case of two time slots

We now develop a pseudo-polynomial dynamic program for the special case of TCTP where $T = 2$ and show how it can be converted into a fully polynomial-time approximation scheme (FPTAS). Together with Corollary 1 this yields that TCTP is weakly NP-hard when $T = 2$. Our approach is inspired by the standard pseudo-polynomial dynamic program for the knapsack problem (Kellerer et al., 2004). An analogous procedure also solves TCSP with $T = 2$, but we omit its description for the sake of brevity.

If $T = 2$, then every feasible schedule is uniquely defined by the set that is tested in the first time slot. For an optimal sequence $\sigma^* = (S^*, N \setminus S^*)$, we refer to the set S^* as an *optimal testing set*. We assume that $n = mT$, which is without loss of generality because one can always add dummy tests with cost 0 and success probability 1 until the instance satisfies this condition. The goal of our dynamic program is to identify such an optimal testing set through a backward recursion based on the index number $i \in \{1, \dots, n\}$ of a job.

For every *stage* $i \in \{1, \dots, n\}$ and *state* $(b, s) \in \{0, \dots, c(N)\} \times \{0, \dots, m\}$, we define the *value function* $v_i(b, s)$ as the minimum joint probability $\prod_{j \in S} p_j$ of a set $S \subseteq \{i, \dots, n\}$ with cost $c(S) = b$ and cardinality $|S| = s$. That is, among the jobs corresponding to the current stage and the higher-indexed stages, we want to select a subset such that, if we schedule these jobs in the first time slot, then the probability that we need to perform the tests in the second time slot is minimal. Here, the state (b, s) specifies the cost and the number of elements that this subset should attain. If b and s are such that no set $S \subseteq \{i, \dots, n\}$ with $c(S) = b$ and $|S| = s$ exists, then we define $v_i(b, s) := +\infty$.

We now discuss how to compute the value function $v_i(b, s)$ for every stage and state recursively. As the boundary condition, let $v_{n+1}(b, s) := 1$ if $b = s = 0$ and $v_{n+1}(b, s) := +\infty$ otherwise. For every stage $i \in \{1, \dots, n\}$ and state $(b, s) \in \{0, \dots, c(N)\} \times \{0, \dots, m\}$, it then holds that

$$v_i(b, s) = \min\{p_i v_{i+1}(b - c_i, s - 1), v_{i+1}(b, s)\} \quad (8)$$

if $v_{i+1}(b - c_i, s - 1)$ is defined and finite, and $v_i(b, s) = v_{i+1}(b, s)$ otherwise. Here, the option $p_i v_{i+1}(b - c_i, s - 1)$ corresponds to selecting test i in that stage and state, whereas the other option refers to not selecting test i . Starting from stage n , Equation (8) allows us to compute the value function $v_1(b, s)$ for every state (b, s) through a backward recursion. By keeping track of the selected tests, we also obtain a testing set that minimizes the joint probability.

Once we know the value function $v_1(b, m)$ for every $b \in \{0, 1, \dots, c(N)\}$, we can solve TCTP with $T = 2$ by enumerating over all these possible costs b and selecting one b^* that minimizes $b +$

$v_1(b, m)(c(N) - b)$. The accompanying set of tests that attains this cost b^* with minimum joint probability $v_1(b^*, m)$ then forms an optimal testing set. Observe that we can indeed limit ourselves to states (b, s) with $s = m$ in the first stage since we know that exactly m jobs need to be scheduled in the first time slot as a consequence of our assumption that $n = mT = 2m$.

Since the computational effort depends on the number $c(N)$, the dynamic program has a pseudo-polynomial running time. By appropriately rounding the costs $(c_j)_{j \in N}$, however, we can convert this pseudo-polynomial dynamic program into a FPTAS. This procedure is analogous to the standard scaling approach for the knapsack problem (Kellerer et al., 2004). Theorem 4, proven in Appendix A, summarizes this result.

Theorem 4. *Given a TCTP instance $(m, n, T, (c_j, p_j)_{j \in N})$ with $T = 2$, let $\sigma^* = (S^*, N \setminus S^*)$ be an optimal testing sequence. For each $\varepsilon > 0$, we can obtain a sequence $\sigma_\varepsilon = (S_\varepsilon, N \setminus S_\varepsilon)$ with $z(\sigma_\varepsilon) \leq (1 + \varepsilon)z(\sigma^*)$ in time $O(n^5/\varepsilon)$.*

Proof. See Appendix A. □

7 Mixed integer programming formulations

In this section, we describe a partial-order-based and an assignment-based mixed integer programming formulation for both TCTP and TCSP. We first address the testing variant in detail and then indicate how the corresponding formulation can be adapted to also solve the search variant. The partial-order-based formulation can be easily modified to also incorporate precedence constraints and, as such, it can be used to solve the precedence-constrained sequential testing problem (Monma and Sidney, 1979). The assignment-based formulation, in turn, can be modified to incorporate a fixed cost for every used time slot and thus to solve the batch-testing problem (Daldal et al., 2016). This illustrates that our approach, based on a linearization of the objective function, is sufficiently general to be useful to solve related testing problems as well.

7.1 Partial-order-based formulation

Inspired by Potts (1980) and Wagner and Davis (2001), our first formulation relies on binary decision variables that indicate the partial order defined by a schedule. Additionally, we include decision variables that keep track of the probability that a certain component (or location) needs to be tested (or searched).

7.1.1 Testing variant

For every two tests $i, j \in N$, define a binary decision variable δ_{ij} that equals 1 if test i is performed before test j and 0 otherwise, and a binary decision variable $\mu_{\{i, j\}}$ that equals 1 if tests i and j are performed simultaneously. The brackets in the subscript of $\mu_{\{i, j\}}$ indicate that the order of indices i and j is of no importance. Next, for every test $i \in N$ and integer $j \in \{0, \dots, n\}$, define a continuous decision variable α_{ij} that equals the probability that test i needs to be performed, conditional on the information that all tests $j + 1, \dots, n$ will be successful. That is, given a set $S \subset N$ of tests performed before test $i \in N$, we want for every $j = 0, \dots, n$ that $\alpha_{ij} = \prod_{l \in S: l \leq j} p_l$ and, therefore, that $\alpha_{in} = \prod_{l \in S} p_l$ equals the probability that test i needs to be performed.

The following mixed integer program then constitutes a valid formulation for TCTP if $n = mT$.

$$\min \sum_{i \in N} c_i \alpha_{in} \tag{9a}$$

$$\begin{aligned}
\text{s.t.} \quad & \delta_{ij} + \delta_{ji} + \mu_{\{i,j\}} = 1 && \forall i, j \in N: i \neq j && (9b) \\
& \mu_{\{i,j\}} + \delta_{ij} + \delta_{jk} - \delta_{ik} \leq 1 && \forall i, j, k \in N: i \neq j \neq k \neq i && (9c) \\
& \sum_{j \in N \setminus \{i\}} \mu_{\{i,j\}} = m - 1 && \forall i \in N && (9d) \\
& \alpha_{i0} = 1 && \forall i \in N && (9e) \\
& \alpha_{ij} \geq \alpha_{i,j-1} - \delta_{ji} && \forall i \in N \text{ and } j \in \{1, \dots, n\} && (9f) \\
& \alpha_{ij} \geq p_j \alpha_{i,j-1} && \forall i \in N \text{ and } j \in \{1, \dots, n\} && (9g) \\
& \delta_{ij}, \mu_{\{i,j\}} \in \{0, 1\} && \forall i, j \in N && (9h)
\end{aligned}$$

Objective function (9a) relates the probability that a test needs to be performed with its cost. Constraints (9b)-(9d) together with the binary constraints enforce that the variables δ_{ij} and $\mu_{\{i,j\}}$ define a partial order of the tests such that in each time slot exactly m tests are performed. Constraints (9e)-(9g) and the non-negative coefficients c_i in the objective function imply that α_{ij} equals the probability that test i needs to be performed conditional on the information that all tests $j + 1, \dots, n$ will be successful, as desired. Indeed, if $\alpha_{i,j-1}$ equals the probability that test i needs to be performed conditional on the information that all tests j, \dots, n will be successful, then α_{ij} equals $p_j \alpha_{i,j-1}$ if test j is performed before test i , and $\alpha_{ij} = \alpha_{i,j-1}$ otherwise.

7.1.2 Search variant

For TCSP we adopt the same definition for δ_{ij} and $\mu_{\{i,j\}}$, but now we directly let α_j reflect the probability that location $j \in N$ needs to be searched. That is, α_j equals the probability that the target has not been found at the time that location j is searched, and thus that the cost c_j needs to be paid. Hence, if we consider the objective function

$$\min \sum_{i \in N} c_i \alpha_i$$

and the constraints

$$\pi_i + \sum_{j \in N \setminus \{i\}} \pi_j (\mu_{\{i,j\}} + \delta_{ij}) = \alpha_i \quad \forall i \in N$$

instead of Objective function (9a) and Constraints (9e)-(9g) in the partial-order-based formulation for TCTP, then we obtain a valid formulation for TCSP. Observe that also for TCSP, the assumption that $n = mT$ is without loss of generality because we can add dummy locations with zero cost and zero probability.

7.1.3 Incorporating precedence constraints

Consider a generalization of TCTP or TCSP in which there are *precedence constraints* defined by a strict partial order $A \subset N \times N$ on N , where $(i, j) \in A$ reflects that component $j \in N$ can only be tested or searched if component $i \in N$ has been tested or searched before. Such constraints can easily be incorporated in our partial-order-based formulations by just adding a constraint $\delta_{ij} = 1$ for each $(i, j) \in A$.

By additionally requiring that $m = 1$ and $T = n$, we obtain a compact mixed integer program for the precedence-constrained sequential testing problem. Rostami et al. (2019) describe a dynamic program and a branch-and-price algorithm for the precedence-constrained sequential testing problem for n -out-of- n systems. Their branch-and-price approach also relies on a

related linear-order-based formulation, but the authors use a Dantzig-Wolfe decomposition with exponentially many decision variables to linearize the objective.

7.2 Assignment-based formulation

In the MIP formulations below, we use binary decision variables that assign the tests (or searches) to the time slots. Additionally, we include decision variables that keep track of the probability that the system's state is still unknown (or that the target is not found) at the beginning of a time slot.

7.2.1 Testing variant

For each test $j \in N$ and time slot $t \in \{1, \dots, T\}$, define a binary decision variable x_{jt} that equals 1 if test j is scheduled in time slot t and 0 otherwise, and a continuous decision variable y_j equal to the probability that test j will actually be performed. Next, for each $j = 0, \dots, n$ and $t \in \{1, \dots, T\}$, define a continuous decision variable z_{jt} equal to the probability that all tests scheduled before time t are successful, conditional on the information that all tests $i > j$ in time slot $t - 1$ are successful. Hence, z_{nt} equals the probability that the tests scheduled in time slot t need to be performed. This leads to the following MIP formulation for TCTP.

$$\min \quad \sum_{j \in N} c_j y_j \quad (10a)$$

$$\text{s.t.} \quad \sum_{t=1}^T x_{jt} = 1 \quad \forall j \in N \quad (10b)$$

$$\sum_{j \in N} x_{jt} \leq m \quad \forall t = 1, \dots, T \quad (10c)$$

$$y_j \geq z_{nt} - 1 + \sum_{s=1}^t x_{js} \quad \forall j \in N \text{ and } t = 1, \dots, T \quad (10d)$$

$$z_{j1} = 1 \quad \forall j = 0, \dots, n \quad (10e)$$

$$z_{0t} = z_{n,t-1} \quad \forall t = 2, \dots, T \quad (10f)$$

$$z_{jt} \geq z_{j-1,t} - x_{j,t-1} \quad \forall j = 1, \dots, n \text{ and } t = 2, \dots, T \quad (10g)$$

$$z_{jt} \geq p_j z_{j-1,t} \quad \forall j = 1, \dots, n \text{ and } t = 2, \dots, T \quad (10h)$$

$$x_{jt} \in \{0, 1\} \quad \forall j \in N \text{ and } t = 1, \dots, T \quad (10i)$$

Objective function (10a) combines the probability that a test needs to be performed with its cost. Constraints (10b)-(10c) together with the binary constraints enforce that every test is scheduled in exactly one time slot, and that at most m tests are scheduled within a time slot. Constraints (10d) together with the objective function and the fact that z_{nt} is non-increasing in t make sure that if test j is scheduled in or before time slot t , it will be performed with a probability of at least z_{nt} . Constraints (10e)-(10h), finally, enforce that the probability z_{jt} equals $p_j z_{j-1,t}$ if test j is scheduled in time slot $t - 1$, and $z_{j-1,t}$ otherwise.

7.2.2 Search variant

For TCSP we adopt the same definition for x_{jt} and y_j , but now we directly let z_t reflect the probability that the locations scheduled in time slot $t \in \{1, \dots, T\}$ need to be searched. That

is, z_t equals the probability that the object is hidden in one of the locations searched at time t or later. Hence, if we consider the constraints

$$y_j \geq z_t - 1 + \sum_{s=1}^t x_{js} \quad \forall j \in N \text{ and } t = 1, \dots, T$$

$$z_t = \sum_{k=t}^T \sum_{j \in N} \pi_j x_{jk} \quad \forall t = 2, \dots, T$$

instead of Constraints (10d)-(10h) in the assignment-based formulation for TCTP, then we obtain a valid formulation for the TCSP.

7.2.3 Incorporating a fixed cost per used time slot

Consider a generalization of TCTP (or TCSP) in which there is a, possibly time-dependent, fixed cost $\beta_t \in \mathbb{R}_{\geq 0}$ to test (or search) a batch of components (or locations) at time $t \in \{1, \dots, T\}$. This cost β_t could reflect a fixed start-up cost, but possibly also a penalty for not knowing the system's state at time t . We can incorporate such fixed costs into the assignment-based formulation for TCTP by introducing for each time slot $t \in \{1, \dots, T\}$ a decision variable u_t that equals the probability that one or more tests will take place in time slot t . We then add to the objective function (10a) an additional summation $\sum_{t=1}^T \beta_t u_t$, and we also include the constraints

$$u_t \geq z_{nt} - 1 + \sum_{j \in N} x_{jt}$$

and $u_t \geq 0$ for every $t \in \{1, \dots, T\}$. An analogous approach allows us to include fixed costs into the assignment-based formulation for TCSP as well.

By setting $m = T = n$ such that there is no deadline or constraint on the number of tests that can be performed in parallel, we obtain a compact exact mixed integer program for the batch-testing problem. Daldal et al. (2016) also describe an assignment-based formulation for this problem, but they approximate the objective function to obtain linearity.

8 Local search heuristic

We now describe a local search procedure for TCTP and TCSP that exploits the fact that, for a given partition, we can determine an optimal schedule efficiently (Lemma 2). In particular, we represent a solution by means of an unordered partition $\{S_1, \dots, S_T\}$ of the set N into T sets with at most m elements each, and translate this partition into a schedule by sequencing the sets in non-decreasing order of their cost-to-probability ratio.

The local search procedure for TCTP works as follows. Starting from a partition, a move consists of either the pairwise interchange between tests in different sets or the insertion of a test in another time slot in which less than m machines are being used. Moreover, after each interchange or insertion, the sets are again sorted in non-decreasing order of their cost-to-probability ratio. If there exists such a move that leads to a strict improvement, then we restart the procedure with this new partition. Otherwise, the partition defines a *locally optimal testing sequence* $\sigma^{\text{ls}} = (S_1^{\text{ls}}, \dots, S_T^{\text{ls}})$ and the procedure ends. The local search for TCSP works in a completely analogous fashion.

As illustrated by the following example, this local search procedure does not provide a constant-factor approximation guarantee for TCTP in general. The locality gap can in par-

ticular be $\Omega(m)$, even if $T = 2$. The example does not seem to be adaptable to TCSP, as in this latter problem all probabilities add up to one (i.e., $\sum_{j \in N} \pi_j = 1$).

Example 2. For given constants $c, k, M \in \mathbb{N}$ all strictly greater than 1, define $p = (c + kM)^{-1/k}$. Consider a TCTP instance with $m = 2k$ machines, deadline $T = 2$, and the following $n = 2k + 1$ tests:

- $I_1 := \{1, \dots, k\}$ with $p_i = p$ and $c_i = c - 1$ for every $i \in I_1$;
- $I_2 := \{k + 1, \dots, 2k\}$ with $p_i = 1$ and $c_i = M$ for every $i \in I_2$;
- $i^* := 2k + 1$ with $p_{i^*} = 0$ and $c_{i^*} = c$.

For given values of c and k , we will derive conditions for the parameter M such that the schedule $\sigma^{\text{ls}} = (I_1, \{i^*\} \cup I_2)$ is locally optimal and the schedule $\sigma^* = (\{i^*\}, I_1 \cup I_2)$ is globally optimal. Since, using $p = (c + kM)^{-1/k}$ and $m = 2k$,

$$z(\sigma^{\text{ls}}) = k(c - 1) + p^k(c + kM) = k(c - 1) + 1 = \frac{m}{2}(c - 1) + 1,$$

whereas $z(\sigma^*) = c$, this would then yield the $\Omega(m)$ locality gap.

Observe that the potentially interesting moves from σ^{ls} are to either interchange a test in I_1 with i^* or to move a test in I_1 to the second time slot. Interchanging a test in I_1 with the test i^* leads to a cost $(k - 1)(c - 1) + c = k(c - 1) + 1 = z(\sigma^{\text{ls}})$ and, as such, it does not yield a strict improvement. Moving a test in I_1 to the second time slot, in turn, leads to a cost $(k - 1)(c - 1) + p^{k-1}(c + c - 1 + kM)$. By substituting the definition of p , we thus obtain that σ^{ls} is a local optimum if

$$(k - 1)(c - 1) + \frac{2c - 1 + kM}{c + kM}(c + kM)^{1/k} \geq k(c - 1) + 1,$$

which is equivalent to requiring that

$$\frac{c^2 + ckM}{2c - 1 + kM}(c + kM)^{-1/k} \leq 1.$$

Since the left-hand side of this latter inequality tends to zero as M increases, there exists a value $M \in \mathbb{N}_0$ such that this inequality is satisfied. For every $k, c \in \mathbb{N}_0$, we can thus choose M sufficiently large such that σ^{ls} is a local optimum, which establishes the $\Omega(m)$ locality gap.

9 Computational experiments

This section reports on the computational performance of our proposed solution methods. After describing the dataset of test instances and the implementation details (Section 9.1), we discuss the performance of our mixed integer programming formulations (Section 9.2) as well as of our local search heuristic (Section 9.3).

9.1 Instance description and implementation details

A set of TCTP and TCSP instances for different values of m and T , with $n = mT$, was randomly generated as follows. For each $j \in N$, an integer cost c_j and weight w_j was drawn randomly, i.e., with uniform probability, from the interval $[0, 10]$ and $[0, 1000]$, respectively. For TCSP, we then let $\pi_j = w_j/w(N)$ for each $j \in N$. For TCTP, in turn, a joint success probability $q = \prod_{j \in N} p_j$

was first drawn randomly from the interval $[0.01, 0.30]$, $[0.31, 0.60]$, or $[0.61, 0.90]$ depending on the experimental setting. This probability was subsequently distributed randomly over the tests by setting $\log(p_j) = \log(q)w_j/w(N)$ for each $j \in N$. By controlling the joint success probability, we avoid that the product $\prod_{j \in N} p_j$ becomes very small as n grows, which could cause numerical issues. For each setting, 10 instances were generated, leading to 30 TCTP and 10 TCSP instances for each (m, T) -combination.

All our algorithms were implemented using the C++ programming language, compiled with Microsoft Visual C++ 14.0, and run using an Intel Core i7-4790 processor with 3.60 GHz CPU speed and 8 GB of RAM under a Windows 10 64-bit OS. To solve the instances, we used the commercial solver IBM ILOG CPLEX 12.8 with a single thread and a time limit of 30 minutes. Lazy cuts were used to implement Constraints (9c) of the partial-order-based formulation, and the solution provided by the local search procedure of Section 8 was used as a warm start. Apart from this, all CPLEX parameters were set to their default values.

To obtain a starting solution for each instance, our local search procedure was run with three different initializations; the resulting solutions were all provided to the solver. For TCTP, the construction of an initial partition was based on either the costs c_j , the probabilities p_j , or the ratios $c_j/(1 - p_j)$, respectively. In particular, all tests were first sorted in non-decreasing value of the corresponding quantity (where those tests $j \in N$ with $p_j = 1$ are put first in this order if $c_j = 0$, and last otherwise), and an initial partition was subsequently obtained by filling the different time slots using the resulting order. For TCSP, an analogous approach based on the costs c_j , probabilities π_j , or ratios c_j/π_j was used.

In order to evaluate the computational performance of our formulations, we use the average CPU time in seconds, the number of solved instances within the 30-minute time limit, the average LP gap, and the average final gap. Here, the LP gap equals the percentage ratio between the value of the optimal solution and the lower bound provided by the linear program obtained by relaxing the integrality constraints in the corresponding MIP formulation. The final gap, in turn, equals the percentage by which the best found solution exceeds the current global lower bound when the 30-minute time limit is reached. The average CPU time and LP gap are computed using solved instances only, whereas the average final gap is computed using unsolved instances only. If all instances for a given (m, T) -combination remained unsolved by a method, then no instances with higher values for m and T were attempted with that method.

9.2 Performance of mixed integer programming formulations

Table 1 displays the computational performance of our TCTP and TCSP formulations. The first and most remarkable aspect emerging from this table is that TCSP appears to be considerably easier to solve than TCTP. While we can solve almost all TCSP instances with $n = mT \leq 40$, regardless of the value for m and T , we have to settle for a much smaller instance size $n \leq 16$ for TCTP. Somewhat surprisingly, we also find that the assignment-based formulation performs better than the partial-order-based formulation for TCTP, whereas this is the other way round for TCSP.

The partial-order-based formulation for TCTP performs quite poorly when both T and m exceed 2. This is somewhat surprising, given that its LP gap is significantly smaller than the one of the assignment-based formulation. One possible explanation is that the partial-order-based formulation, with $O(n^3)$ constraints, is relatively large. This makes the formulation stronger, but it also increases the time needed to solve the LP relaxation in every node of the branch-and-bound tree. For both formulations, the performance seems to be more sensitive to the value of the deadline T than to the number m of testers. In fact, while all instances with $m = 4$ and $T = 4$ can be solved to optimality, there is only one solved instance with $m = 2$ and $T = 8$.

Table 1: Computational performance of partial-order-based and assignment-based formulations for TCTP and TCSP. The averages for the CPU time (in seconds) and percentage LP gap (%LP) are computed using solved instances only (# out of 30 for TCTP and out of 10 for TCSP, with a 30-minute time limit). The average final gap (%fn) is computed using unsolved instances only.

m	T	Time-critical testing problem								Time-critical search problem							
		Partial order				Assignment				Partial order				Assignment			
		CPU	#	%LP	%fn	CPU	#	%LP	%fn	CPU	#	%LP	%fn	CPU	#	%LP	%fn
2	2	0.1	30	27.08		0.1	30	24.45		0.1	10	0.06		0.0	10	17.22	
2	3	0.1	30	26.68		0.1	30	55.48		0.1	10	0.63		0.1	10	38.80	
2	4	0.2	30	25.77		0.2	30	67.01		0.1	10	0.56		0.1	10	51.24	
2	5	14.7	30	24.02		4.5	30	76.34		0.1	10	0.45		0.4	10	60.45	
2	6	113.9	29	30.30	0.20	152.1	30	80.47		0.1	10	0.50		1.2	10	66.06	
2	7	439.3	13	26.76	1.55	843.3	7	96.72	18.70	0.1	10	0.49		20.7	10	71.32	
2	8		0		6.10	777.5	1	87.44	44.37	0.1	10	0.29		514.5	8	75.35	8.30
2	9						0		62.04	0.1	10	0.24		207.2	4	77.98	15.48
2	10									0.1	10	0.27		261.6	1	81.93	22.28
4	2	0.3	30	36.68		0.0	30	28.46		0.0	10	0.18		0.0	10	19.44	
4	3	271.1	23	35.20	1.11	0.5	30	52.62		0.1	10	0.47		0.1	10	39.75	
4	4		0		5.50	237.3	30	69.04		0.2	10	1.05		0.7	10	51.05	
4	5						0		31.12	0.4	10	0.54		349.1	9	60.57	3.68
4	6									1.0	10	0.43			0		12.80
4	7									3.1	10	0.39					
4	8									140.9	10	0.22					
4	9									224.7	10	0.59					
4	10									213.0	10	0.35					
6	2	277.4	27	36.41	0.47	0.1	30	32.39		0.1	10	0.01		0.1	10	21.09	
6	3		0		11.41	28.8	30	55.26		0.2	10	0.29		0.3	10	39.46	
6	4						0		24.19	0.8	10	0.39		61.5	10	52.42	
6	5									3.4	10	0.38		883.9	1	61.27	11.46
6	6									140.8	10	0.37			0		30.32
6	7									107.4	5	0.27	0.30				
6	8									157.6	5	0.19	0.21				
6	9									220.1	2	0.16	0.24				
6	10									0		0.27					
8	2	921.4	7	72.61	4.02	0.1	30	29.45		0.1	10	0.00		0.1	10	16.60	
8	3		0		14.44	574.7	25	61.52	1.94	0.6	10	0.26		0.6	10	36.87	
8	4						0		38.08	8.0	10	0.29		637.9	3	52.43	5.93
8	5									66.6	10	0.31			0		23.80
8	6									448.6	7	0.27	0.17				
8	7									458.1	6	0.17	0.30				
8	8									108.4	1	0.09	0.15				
8	9									0		0.34					
10	2		0		13.96	0.1	30	33.39		0.1	10	0.02		0.1	10	19.60	
10	3					1,210.3	1	60.97	13.10	1.9	10	0.26		32.9	10	37.77	
10	4						0		43.30	22.6	10	0.25			0		9.56
10	5									284.7	10	0.20					
10	6									926.9	5	0.25	0.17				
10	7									0		0.40					

Table 2: Computational performance for the partial-order based formulation for larger TCSP instances.

(a) Partial-order based formulation for TCSP with larger values for the deadline T .

T	$m = 2$				$m = 4$			
	CPU	#	%LP	%fin	CPU	#	%LP	%fin
15	0.5	10	0.20		1,532.9	1	0.12	0.29
20	25.6	10	0.14			0		0.26
25	48.8	10	0.09					
30	173.1	10	0.06					
35	695.5	5	0.05	0.06				
40	99.7	1	0.03	0.05				
45	890.2	2	0.02	0.08				
50		0		0.09				

(b) Partial-order based formulation for TCSP with larger values for the number of searchers m .

m	$T = 2$				$T = 3$				$T = 4$				$T = 5$			
	CPU	#	%LP	%fin	CPU	#	%LP	%fin	CPU	#	%LP	%fin	CPU	#	%LP	%fin
12	0.5	10	0.01		4.8	10	0.19		246.1	10	0.31		113.3	4	0.09	0.26
14	0.6	10	0.02		6.1	10	0.06		370.2	10	0.19		612.7	4	0.04	0.32
16	0.4	10	0.00		12.1	10	0.05		472.6	7	0.09	0.28	1,280.4	2	0.04	0.53
18	2.0	10	0.01		26.0	10	0.01		713.2	5	0.07	0.12	205.2	2	0.00	0.72
20	1.5	10	0.00		97.1	10	0.05		757.0	7	0.02	0.16		0		0.78
22	1.8	10	0.00		246.4	10	0.07		1,015.8	2	0.09	0.70				
24	3.7	10	0.03		457.3	10	0.08		567.2	2	0.02	0.45				
26	9.3	10	0.00		395.8	5	0.04	0.31		0		0.86				
28	18.9	10	0.00		497.8	5	0.02	0.53		0		0.85				
30	8.4	10	0.00		1,184.6	7	0.01	0.89		0		0.89				

For TCSP, in contrast, the partial-order-based formulation clearly outperforms the assignment-based formulation. This superior performance is most likely explained by the remarkably small LP gap, which is several orders of magnitude below the gap of the assignment-based formulation. As such, our results are consistent with the earlier finding that partial-order-based formulations typically perform well for scheduling problems with a total weighted completion time objective function (see for instance Potts, 1980 or Keha et al., 2009).

Table 2 reports on the computational performance for the partial-order-based formulation for larger TCSP instances. It shows that the time limit appears to become binding only when $n = mT$ reaches 60. When comparing Tables 2a and 2b, we observe that, similarly to TCTP, instances with low T and high m seem to be easier to solve than instances with high T and low m in general. For the special case where $m = 2$, however, the partial-order-based formulation can still solve instances with a fairly high value for T . Observe in this respect that, although we have established in Section 5 that m -TCSP is strongly NP-hard if $m \geq 3$, the complexity status of TCSP in the special case where $m = 2$ is still open.

Table 3 further details the computational performance for TCTP by disaggregating the results on the basis of the success probability q . Here, instances with lower values of q seem to be relatively more tractable. To explain this behavior, recall that a smaller q results in a wider range for the values that the objective function can attain across the various solutions (Lemma 3). This wider range might then be exploited by the solver to differentiate more easily among different solutions. The effect of q on the LP gap is also consistent with Lemma 3.

For $T = 2$, the dynamic program described in Section 6 provides an alternative approach to obtain an optimal schedule. To test the computational performance of this dynamic program, we have performed additional experiments using the same set of instances as described above, except that we have re-generated the costs from intervals of increasing sizes (instead of the $[0, 10]$ interval as used previously) to examine how this affects performance. Table 4 reports on

Table 3: Computational performance of the TCTP formulations for different values of $q = \prod_{j \in N} p_j$.

m	T	Partial-order-based formulation									Assignment-based formulation								
		$q \in [0.01, 0.30]$			$q \in [0.31, 0.60]$			$q \in [0.61, 0.90]$			$q \in [0.01, 0.30]$			$q \in [0.31, 0.60]$			$q \in [0.61, 0.90]$		
		CPU	#	%LP	CPU	#	%LP	CPU	#	%LP	CPU	#	%LP	CPU	#	%LP	CPU	#	%LP
2	2	0.1	10	49.31	0.1	10	23.96	0.1	10	7.95	0.1	10	27.59	0.1	10	32.04	0.1	10	13.73
2	3	0.1	10	48.84	0.1	10	23.63	0.1	10	7.58	0.1	10	69.31	0.1	10	65.37	0.1	10	31.77
2	4	0.3	10	48.58	0.2	10	22.19	0.1	10	6.54	0.1	10	76.59	0.2	10	79.36	0.2	10	45.08
2	5	12.3	10	46.04	10.5	10	19.03	21.3	10	6.98	1.9	10	86.93	5.3	10	86.55	6.3	10	55.53
2	6	103.7	10	61.86	78.6	10	20.38	164.5	9	6.24	18.1	10	87.51	155.9	10	92.25	282.4	10	61.64
2	7	754.6	5	50.86	312.9	5	15.05	124.6	3	6.09	843.3	7	96.72						
2	8										777.5	1	87.44						
4	2	0.3	10	68.49	0.2	10	29.67	0.3	10	11.88	0.0	10	35.98	0.0	10	34.35	0.0	10	15.05
4	3	139.9	10	60.31	308.4	6	23.19	426.5	7	9.65	0.4	10	67.25	0.5	10	61.49	0.7	10	29.13
4	4										94.4	10	78.99	217.8	10	79.94	399.9	10	48.21
6	2	49.9	10	56.96	295.7	9	33.40	541.0	8	14.09	0.1	10	45.78	0.1	10	35.94	0.1	10	15.45
6	3										13.9	10	67.44	27.5	10	65.65	44.9	10	32.68
8	2	789.4	6	78.49	1,713.0	1	37.33				0.1	10	40.08	0.1	10	32.33	0.1	10	15.94
8	3										417.1	10	71.63	602.7	9	65.86	795.5	6	38.15
10	2										0.1	10	47.05	0.1	10	39.88	0.1	10	13.22
10	3										1,210.3	1	60.97						

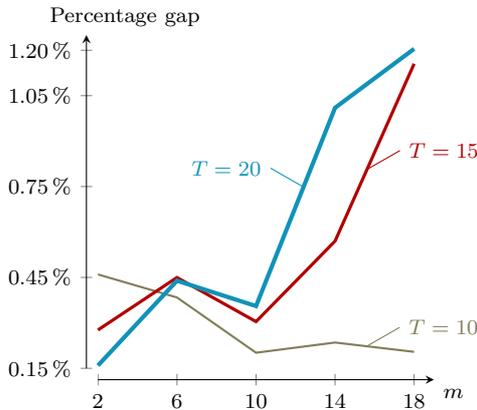
Table 4: CPU time in seconds for the assignment-based formulation and the dynamic program (DP) of Section 6, for TCTP and TCSP instances with $T = 2$ and $m = 10$, depending on the range from which the costs c_j are drawn.

Interval for c_j	TCTP		TCSP	
	Assignment	DP	Assignment	DP
[0, 10]	0.0634	0.0004	0.0381	0.0003
[0, 100]	0.0652	0.0027	0.0408	0.0023
[0, 1 000]	0.0621	0.0291	0.0439	0.0255
[0, 10 000]	0.0757	0.2454	0.0394	0.2572
[0, 100 000]	0.0763	2.8271	0.0525	3.0296

Table 5: Computational performance of our local search for TCTP and TCSP. The average and maximal CPU time (both in seconds) are computed over all instances for which an optimal solution is known. The column %opt gives the percentage of instances for which the local search found the optimum, and the average and maximal percentage optimality gap is computed over instances for which the local search did not find the optimal solution.

Problem	avg CPU	max CPU	%opt	avg %gap	max %gap
TCTP	0.026	0.064	100.00	0.000	0.000
TCSP	0.039	0.070	96.36	0.050	0.131

Figure 1: Average percentage gap between the upper bound provided by the local search procedure and the lower bound provided by the LP relaxation of our formulations, for large TCSP instances.



the results of these additional experiments. We focus on the assignment-based formulation and on instances with $m = 10$ because, for this setting, we can solve all instances for both TCTP and TCSP. The table clearly shows that the required CPU time for the dynamic program is roughly proportional to the interval size, which is not surprising given the theoretical pseudo-polynomial running time of $O(n^2c(N))$. The assignment-based formulation, in contrast, seems to be mostly unaffected by the interval size, and it outperforms the dynamic program for intervals of sizes 10 000 and 100 000.

9.3 Performance of the local search procedure

Table 5 indicates that the local search described in Section 8, with the three initializations as discussed in Section 9.1, finds near-optimal solutions in very limited computation times. In fact, the local search is able to find the optimal solution for all TCTP instances solved to optimality by the MIP formulations. For TCSP, it finds the optimum for 96.47% of the solved instances and, for those instances where it does not find the optimum, the maximal optimality gap is only 0.131%.

To evaluate the performance of our local search for large TCSP instances, Figure 1 shows the average percentage gap between the upper bound provided by the local search procedure and the lower bound provided by the LP relaxation of our formulations. Since the LP relaxation for the partial-order-based formulation could not always be solved to optimality within the 30-minute time limit, we added Constraints (9c) iteratively in a cutting-plane fashion; the considered lower

bound is then the best bound obtained within the time limit. The gap displayed in Figure 1 thus overestimates the true locality gap as it is based on a lower bound for the optimal search cost. As even this overestimate never exceeds 1.2%, one can infer that our local search provides high-quality solutions for large TCSP instances as well. Here, the relatively larger gaps for $T = 15$ and $T = 20$ when $m \geq 14$ can be partially explained by the fact that we could not solve the LP relaxation to optimality for such large instances. We do not report a similar analysis for TCTP, as our TCTP formulations are significantly weaker than the TCSP formulations, and the lower bound provided by their LP relaxations is not sufficiently tight to assess the performance of the local search.

10 Conclusion

We have introduced the time-critical testing problem and the time-critical search problem. These natural generalizations of the classical sequential testing problem and search problem are applicable in a wide range of time-critical operations, such as machine maintenance, medical diagnosis, and new product development. We have shown that both problems are NP-hard, we have developed a pseudo-polynomial dynamic program as well as a FPTAS for the special case of two time slots, and we have described two mixed integer programming formulations as well as a local search heuristic for the general case. With minor modifications, these formulations can also solve closely related testing problems for which no exact and compact mixed integer programming formulation was previously available. Based on extensive computational experiments, we find that our assignment-based formulation performs better than our partial-order-based formulation for the testing variant, while this is the other way round for the search variant. Despite its lack of a theoretical worst-case guarantee, we also find that, empirically, our local search procedure performs very well in finding near-optimal solutions in limited computation times.

Although we have established that both problems are strongly NP-hard for every fixed number of machines $m \geq 3$, their complexity status is still open for $m = 2$. An interesting open question is thus whether there exists a polynomial-time algorithm for the special case of two machines, and whether the search and testing variant still have the same complexity status for this special case. A second open question is whether TCTP is weakly or strongly NP-hard in the special case of a fixed deadline $T \geq 3$. Corollary 1 combined with the pseudo-polynomial dynamic program of Section 6 imply that the problem is weakly NP-hard for $T = 2$, but it does not seem straightforward to generalize the dynamic program to arbitrary but fixed T .

Another promising direction for further research is to develop approximation algorithms for both problem variants. For the testing variant, we have illustrated that a greedy approach that iteratively tests a subset with minimal cost-to-failure-probability as well as a pairwise-interchange-based local search heuristic fail to provide a constant-factor approximation guarantee. The approaches of Daldal et al. (2016) and Segev and Shaposhnik (2018) for the batch testing problem also do not seem to be directly modifiable towards our setting. For the search variant, in turn, an analogous cost-to-probability-based greedy approach also does not provide a constant-factor approximation guarantee, but it is still an open question whether our local search heuristic provides a constant-factor approximation for this variant.

References

Agnētis, A., Detti, P. and Pranzo, M. (2014), ‘The list scheduling algorithm for scheduling unreliable jobs on two parallel machines’, *Discrete Applied Mathematics* **165**, 2–11.

- Agnetis, A., Detti, P., Pranzo, M. and Sodhi, M. S. (2009), ‘Sequencing unreliable jobs on parallel machines’, *Journal of Scheduling* **12**(1), 45.
- Agnetis, A. and Lidbetter, T. (2020), ‘The largest-Z-ratio-first algorithm is 0.8531-approximate for scheduling unreliable jobs on m parallel machines’, *Operations Research Letters* **48**(4), 405–409.
- Alpern, S., Fokkink, R., Gasieniec, L., Lindelauf, R. and Subrahmanian, V. (2013), *Search Theory*, Springer.
- Berend, D., Brafman, R., Cohen, S., Shimony, S. E. and Zucker, S. (2014), ‘Optimal ordering of independent tests with precedence constraints’, *Discrete Applied Mathematics* **162**, 115–127.
- Boothroyd, H. (1960), ‘Least-cost testing sequence’, *Journal of the Operational Research Society* **11**(3), 137–138.
- Daldal, R., Gamzu, I., Segev, D. and Ünlüyurt, T. (2016), ‘Approximation algorithms for sequential batch-testing of series systems’, *Naval Research Logistics* **63**(4), 275–286.
- Daldal, R., Özlük, Ö., Selçuk, B., Shahmoradi, Z. and Ünlüyurt, T. (2017), ‘Sequential testing in batches’, *Annals of Operations Research* **253**(1), 97–116.
- De Reyck, B. and Leus, R. (2008), ‘R&D project scheduling when activities may fail’, *IIE Transactions* **40**(4), 367–384.
- Fokkink, R., Lidbetter, T. and Végh, L. A. (2019), ‘On submodular search and machine scheduling’, *Mathematics of Operations Research* **44**(4), 1431–1449.
- Garey, M. R. (1973), ‘Optimal task sequencing with precedence constraints’, *Discrete Mathematics* **4**(1), 37–56.
- Garey, M. R. and Johnson, D. S. (1979), *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W.H. Freeman and Company.
- Glazebrook, K. and Washburn, A. (2004), ‘Shoot-look-shoot: A review and extension’, *Operations Research* **52**(3), 454–463.
- Gluss, B. (1959), ‘An optimum policy for detecting a fault in a complex system’, *Operations Research* **7**(4), 468–477.
- Gould, J. (1984), ‘On efficient salvo policies’, *Naval Research Logistics* **31**(1), 159–162.
- Happach, F., Hellerstein, L. and Lidbetter, T. (2020), ‘A general framework for approximating min sum ordering problems’, *arXiv preprint arXiv:2004.05954* .
- Hermans, B., Hamers, H., Leus, R. and Lindelauf, R. (2019), ‘Timely exposure of a secret project: Which activities to monitor?’, *Naval Research Logistics* **66**(6), 451–468.
- Johnson, S. M. (1956), Optimal sequential testing, Technical report, RAND Corporation.
- Kadane, J. B. (1969), ‘Quiz show problems’, *Journal of Mathematical Analysis and Applications* **27**(3), 609–623.
- Keha, A. B., Khowala, K. and Fowler, J. W. (2009), ‘Mixed integer programming formulations for single machine scheduling problems’, *Computers & Industrial Engineering* **56**(1), 357–367.
- Kellerer, H., Pferschy, U. and Pisinger, D. (2004), *Knapsack Problems*, Springer.
- Kelly, F. (1982), ‘A remark on search and sequencing problems’, *Mathematics of Operations Research* **7**(1), 154–157.
- Lawler, E. L. (1978), Sequencing jobs to minimize total weighted completion time subject to precedence constraints, in B. Alspach, P. Hell and D. Miller, eds, ‘Algorithmic Aspects of Combinatorics’, Vol. 2 of *Annals of Discrete Mathematics*, Elsevier, pp. 75–90.
- Lenstra, J. K. and Rinnooy Kan, A. (1978), ‘Complexity of scheduling under precedence constraints’, *Operations Research* **26**(1), 22–35.

- Li, S. and Huang, S. (2018), ‘Multiple searchers searching for a randomly distributed immobile target on a unit network’, *Networks* **71**(1), 60–80.
- Lidbetter, T. (2020), ‘Search and rescue in the face of uncertain threats’, *European Journal of Operational Research* **285**(3), 1153 – 1160.
- Lin, K. Y. and Singham, D. I. (2016), ‘Finding a hider by an unknown deadline’, *Operations Research Letters* **44**(1), 25–32.
- Mitten, L. (1960), ‘An analytic solution to the least cost testing sequence problem’, *Journal of Industrial Engineering* **11**(1), 17.
- Monma, C. L. and Sidney, J. B. (1979), ‘Sequencing with series-parallel precedence constraints’, *Mathematics of Operations Research* **4**(3), 215–224.
- Potts, C. N. (1980), ‘An algorithm for the single machine sequencing problem with precedence constraints’, *Mathematical Programming Studies* **13**, 78–87.
- Rostami, S., Creemers, S., Wei, W. and Leus, R. (2019), ‘Sequential testing of n -out-of- n systems: Precedence theorems and exact methods’, *European Journal of Operational Research* **274**(3), 876–885.
- Segev, D. and Shaposhnik, Y. (2018), A polynomial-time approximation scheme for sequential batch testing of series systems. Available at SSRN: <https://ssrn.com/abstract=3277805> or <http://dx.doi.org/10.2139/ssrn.3277805>.
- Sidney, J. B. (1975), ‘Decomposition algorithms for single-machine sequencing with precedence relations and deferral costs’, *Operations Research* **23**(2), 283–298.
- Smith, W. E. (1956), ‘Various optimizers for single-stage production’, *Naval Research Logistics Quarterly* **3**(1-2), 59–66.
- Stone, L. D. (2007), *Theory of Optimal Search, 2nd edition*, INFORMS.
- Ünlüyurt, T. (2004), ‘Sequential testing of complex systems: A review’, *Discrete Applied Mathematics* **142**(1-3), 189–205.
- van Ee, M. (2020), ‘On efficient algorithms for finding efficient salvo policies’, *Naval Research Logistics* **67**(2), 147–158.
- Wagner, B. J. and Davis, D. J. (2001), ‘Discrete sequential search with group activities’, *Decision Sciences* **32**(4), 557–574.

Appendix A A FPTAS for TCTP with two time slots

In this appendix, we prove Theorem 4 by showing how the pseudo-polynomial dynamic program of Section 6 can be converted into a FPTAS by appropriately rounding the testing costs. Without loss of generality, assume that the tests are indexed in non-increasing value of their costs, with the largest probability as tie-breaker. That is, $c_1 \geq \dots \geq c_n$, and $p_i \geq p_{i+1}$ for every $i \in N$ with $c_i = c_{i+1}$. If $c(S^*) = 0$, then an optimal schedule is given by $\sigma = (S, N \setminus S)$ with $S = \{j, \dots, n\}$ such that j is the smallest integer for which $c_j = 0$ and $n - j + 1 \leq m$. Since we can therefore check in polynomial time whether $c(S^*) = 0$, we henceforth assume that $c(S^*) > 0$.

Consider an arbitrary $\varepsilon > 0$ and call n_+ the largest index for which $c_{n_+} > 0$. For each $i = 1, \dots, n_+$, we consider a modified instance in which the costs are rounded down to the nearest multiple of $\mu^{(i)} := \varepsilon c_i / n$. That is, we replace the cost of each test $j \in N$ by a new cost $c_j^{(i)} := \lfloor c_j / \mu^{(i)} \rfloor$ and perform our dynamic program for stages $j \in \{i, \dots, n\}$ with the modified costs to compute $v_i(b, m)$ for each $b = 0, \dots, \sum_{j=i}^n c_j^{(i)}$. If $v_i(b, m) = +\infty$ for all values of b , then we define $z_\varepsilon^{(i)} := +\infty$. Otherwise, we consider an optimal testing set $S_\varepsilon^{(i)} \subseteq \{i, \dots, n\}$ and the schedule $\sigma_\varepsilon^{(i)} = (S_\varepsilon^{(i)}, N \setminus S_\varepsilon^{(i)})$ for the modified instance, and we denote the expected testing cost in the original instance as

$$z_\varepsilon^{(i)} := z(\sigma_\varepsilon^{(i)}) = \sum_{j \in S_\varepsilon^{(i)}} c_j + \sum_{j \in N \setminus S_\varepsilon^{(i)}} c_j \prod_{j \in S_\varepsilon^{(i)}} p_j.$$

Observe that $z_\varepsilon^{(i)}$ is finite for at least one $i \in \{1, \dots, n\}$ since $v_1(b, m) \leq 1$ for $b = \sum_{j=1}^m c_j^{(1)}$. The algorithm completes by returning a testing set $S_\varepsilon = S_\varepsilon^{(k)}$ and schedule $\sigma_\varepsilon = (S_\varepsilon, N \setminus S_\varepsilon)$, where k is such that $\sigma_\varepsilon^{(k)}$ achieves the minimum cost $z_\varepsilon^{(k)} = \min_{i=1, \dots, n_+} z_\varepsilon^{(i)}$ for the original instance.

Since $c_1 \leq \dots \leq c_n$, we have for each $i = 1, \dots, n_+$ that

$$\sum_{j=i}^n c_j^{(i)} = \sum_{j=i}^n \left\lfloor \frac{c_j}{\mu^{(i)}} \right\rfloor = \sum_{j=i}^n \left\lfloor \frac{nc_j}{\varepsilon c_i} \right\rfloor \leq \sum_{j=i}^n \frac{nc_j}{\varepsilon c_i} \leq \sum_{j=i}^n \frac{n}{\varepsilon} \leq \frac{n^2}{\varepsilon},$$

and our dynamic program yields $S_\varepsilon^{(i)}$ for a given i in time $O(n^4/\varepsilon)$. Since we need to evaluate $z_\varepsilon^{(i)}$ for at most n different values of i , we thus obtain S_ε and σ_ε in time $O(n^5/\varepsilon)$.

It remains to be shown that $z(\sigma_\varepsilon) \leq (1 + \varepsilon)z(S^*)$. Let i be the largest index for which $S^* \subseteq \{i, i+1, \dots, n\}$. Observe that $z_\varepsilon^{(i)}$ is finite since S^* is feasible, and that $z(S^*) \geq c(S^*) \geq c_i > 0$ because $i \in S^*$ and $c(S^*) > 0$. It also follows from our definition of $S_\varepsilon^{(i)}$ as an optimal testing set in the modified instance that

$$\sum_{j \in S_\varepsilon^{(i)}} c_j^{(i)} + \sum_{j \in N \setminus S_\varepsilon^{(i)}} c_j^{(i)} \prod_{j \in S_\varepsilon^{(i)}} p_j \leq \sum_{j \in S^*} c_j^{(i)} + \sum_{j \in N \setminus S^*} c_j^{(i)} \prod_{j \in S^*} p_j.$$

Finally, observe that, by construction, $c_j^{(i)} = \lfloor c_j / \mu^{(i)} \rfloor \leq c_j / \mu^{(i)} \leq \lfloor c_j / \mu^{(i)} \rfloor + 1$ for every $j \in N$. From the observations above and the definition of $\mu^{(i)} = \varepsilon c_i / n$, we obtain that

$$\begin{aligned} z(\sigma_\varepsilon^{(i)}) &= \sum_{j \in S_\varepsilon^{(i)}} c_j + \sum_{j \in N \setminus S_\varepsilon^{(i)}} c_j \prod_{j \in S_\varepsilon^{(i)}} p_j \\ &\leq \sum_{j \in S_\varepsilon^{(i)}} \mu^{(i)} \left(\left\lfloor \frac{c_j}{\mu^{(i)}} \right\rfloor + 1 \right) + \sum_{j \in N \setminus S_\varepsilon^{(i)}} \mu^{(i)} \left(\left\lfloor \frac{c_j}{\mu^{(i)}} \right\rfloor + 1 \right) \prod_{j \in S_\varepsilon^{(i)}} p_j \end{aligned}$$

$$\begin{aligned}
&\leq \mu^{(i)} \left(\sum_{j \in S_\varepsilon^{(i)}} \left\lfloor \frac{c_j}{\mu^{(i)}} \right\rfloor + \sum_{j \in N \setminus S_\varepsilon^{(i)}} \left\lfloor \frac{c_j}{\mu^{(i)}} \right\rfloor \prod_{j \in S_\varepsilon^{(i)}} p_j \right) + n\mu^{(i)} \\
&\leq \mu^{(i)} \left(\sum_{j \in S^*} \left\lfloor \frac{c_j}{\mu^{(i)}} \right\rfloor + \sum_{j \in N \setminus S^*} \left\lfloor \frac{c_j}{\mu^{(i)}} \right\rfloor \prod_{j \in S^*} p_j \right) + n\mu^{(i)} \\
&\leq \left(\sum_{j \in S^*} c_j + \sum_{j \in N \setminus S^*} c_j \prod_{j \in S^*} p_j \right) + \varepsilon c_i \\
&\leq z(\sigma^*) + \varepsilon z(\sigma^*) = (1 + \varepsilon)z(\sigma^*).
\end{aligned}$$

Since $z(\sigma_\varepsilon) = \min_{j=1, \dots, n_+} z(\sigma_\varepsilon^{(j)})$, this shows that $z(\sigma_\varepsilon) \leq z(\sigma_\varepsilon^{(i)}) \leq (1 + \varepsilon)z(\sigma^*)$. □