Active Diagnosis with Costs and Rewards

Serge Haddad

Université Paris-Saclay, LMF, ENS Paris-Saclay, CNRS, Gif-sur-Yvette, France

Engel Lefaucheux

Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

Stefan Schwoon

Université Paris-Saclay, LMF, ENS Paris-Saclay, CNRS, Gif-sur-Yvette, France

Abstract

Diagnosis is the task of detecting fault occurrences in a partially observed system. Depending on the possible observations, a discrete-event system may be diagnosable or not. Active diagnosis aims at controlling the system to render it diagnosable. In the past, the main analyzed criterion of the quality of an active diagnoser has been the delay between the fault occurrence and its detection. Here we generalize this study by (1) associating costs or rewards with faulty runs, (2) defining three related decision problems, and (3) analyzing their decidability/complexity in the non-deterministic and probabilistic frameworks under several hypotheses. We study non-deterministic and probabilistic semantics and compare their decidability and complexity. In particular, we exhibit one problem decidable for non-deterministic systems but undecidable for probabilistic ones. Furthermore we establish tight lower and upper bounds for the size of the active diagnoser (when it exists).

2012 ACM Subject Classification Mathematics of computing → Stochastic processes

Keywords and phrases Partial observation, diagnosis, game and automata theory, controller synthesis, probabilistic discrete event systems

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2026.37

Related Version *Full Version*: https://inria.hal.science/hal-05667567

Funding This work is partially supported by ANR BisoUS (ANR-22-CE48-0012) and ANR GUMMIS (ANR-25-CE48-5096).

1 Introduction

Diagnosis. Most discrete-event systems are prone to faults independently of the robustness of their design. For instance, faults can be triggered by a pathological behaviour of the environment or the breakdown of some material component. However, when a fault occurs, from the point of view of an external observer the behaviour of the system may seem similar to a correct one. Thus *diagnosis*, which consists in detecting whether the system is surely faulty, is a central task in managing such systems.

This problem can be formalized as a (finite) labeled non-deterministic transition system whose actions are observable and where states are partitioned into correct and faulty ones (with no transition from a faulty to a correct state). A *diagnoser* observes the sequence of actions. When the system enters a faulty state, then it must announce, after a finite number of further actions, that the system is surely faulty. A system is *diagnosable* iff there exists some diagnoser for it. As the available observations may be insufficient to perform diagnosis, *diagnosability* verification has received considerable attention since the seminal paper by Sampath et al [7]; see also [3, 2]. These works construct a dedicated deterministic version of the original plant, such that the absence of a particular kind of cycles in this auxiliary automaton is equivalent to diagnosability. Via a different approach called *twin-plant construction* the diagnosability problem is shown to be solvable in polynomial time [8].

37th International Conference on Concurrency Theory (CONCUR 2026).
Editors: Ana Sokolova and Patrick Totzke; Article No. 37; pp. 37:1–37:16
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Active diagnosis. A useful way to address the diagnosability problem is a two-stage method: first design the system $\mathcal{A}$ in order to ensure its functionalities as long as the system is correct and then design a controller that, while letting the system live, forces the behavior of $\mathcal{A}$ to stay within a diagnosable subset of its behaviors. Sampath et al [6] have initiated such an approach called *active diagnosis*. The pair consisting of the controller and the diagnoser is called an *active diagnoser*. Later, Chanthery and Pencolé [4] have proposed a planning-based approach via a twin plant construction. In [5], the authors show that the active-diagnosability problem (i.e., whether there exists an active diagnoser) is EXPTIME-complete. Moreover in the positive case, the size of the associated active diagnoser returned by their procedure is singly exponential with respect to the size of the system, while that of [6] is doubly exponential, and this size is optimal.

In [1], the authors introduce a probabilistic framework for diagnosability by equipping transitions with positive weights. Once a controller is designed, the weights of allowed transitions outgoing from a state are normalized to probabilities yielding a Markov chain (finite when the controller has a finite memory). In this context, the active diagnoser has to detect that the system is surely faulty with probability one (a weaker requirement). They established that the active-diagnosability problem is EXPTIME-complete while the safe active-diagnosability problem, which additionally requires that the active diagnoser ensures a positive probability to stay correct, is undecidable.

Delays in active diagnosis. In [5], the authors introduce the *delay* of an active diagnoser as the maximal number of actions between the fault occurrence and its detection over all faulty runs of the controlled system. They present a family of systems parametrized by n, the number of states, for which a minimal-delay active diagnoser must have $2^{\Omega(n\log(n))}$ states. For an actively diagnosable system of size n, they exhibit a construction yielding an active diagnoser with size $2^{\mathcal{O}(n)}$, whose delay is at most twice the minimal possible delay.

Costs and rewards in active diagnosis. While the maximal delay from the fault occurrence to its detection is a natural way to measure the quality of the active diagnoser, it has some limitations. First, the actions that occur during the diagnosis process are not equally dangerous. Let us consider the point of view of the company that uses the system: an error while reading its database is significantly less damaging than a faulty update on it. Thus the company aims at minimizing the cumulative *cost* up to the fault detection.

On the other hand, if fault detection comes with a manifest malfunction that prevents further operation, the diagnoser may want to maximize the remaining utility of the system between fault occurrence and its detection. Delaying the malfunction may be possible for a finite or infinite amount of time. In the first case, we speak of maximizing the *reward*. In the latter case, avoiding a possible malfunction may require the system to operate in a degraded mode indefinitely, so that one still wants to perform fault detection eventually, for instance after having put in place contingency measures. We call this the problem of active diagnosis with *unbounded rewards*.

Our contributions. We introduce the model of initialized weighted controllable systems (IWCS) with correct and faulty states. Transitions between faulty states are equipped with integer weights that are accumulated along a faulty run. Then we state three decision problems: Given an IWCS and some threshold B, the *cost problem* (resp. *reward problem*) asks whether there exists an active diagnoser such that the weight of all faulty runs is at most (resp. at least) B. Given an IWCS, the *unbounded reward problem* asks whether for every threshold B there is an active diagnoser such that the weight of all faulty runs is at least B. There are different ways to restrict IWCS: (1) the weights are observable, i.e. they only depend on the actions and not on the transitions (called weight-observable IWCS) (2) the weights are non negative/positive (called non negative/positive IWCS) (3) there is a single action labelling the transitions from correct states to faulty ones, and this action labels only transitions outgoing from correct states (called SF-IWCS) and (4) this action only occurs initially (called SIF-IWCS). We show that when the weights are not observable these problems are undecidable even for SIF-IWCS. Thus we establish several incomparable restrictions that ensure the decidability:

- the three problems for weight-observable SF-IWCS;
- the cost and reward problems for non negative IWCS;
- the unbounded reward problem for positive IWCS;
- the unbounded reward problem for non negative weight-observable IWCS.

Furthermore we establish tight lower and upper bounds for the required memory size of an active diagnoser achieving some threshold w.r.t. the cost or the reward.

We then introduce the model of initialized weighted controllable probabilistic transition systems (pIWCS), where the choice between the enabled transitions is randomly performed. Given some goal (as usual in a probabilistic framework), we only require that the active diagnoser almost surely achieves it. We prove that, contrary to the non deterministic case, the three problems are undecidable even for weight-observable SIF-IWCS. However we show that the cost and reward problems are decidable for non negative pIWCS and moreover with the same computational complexity as the one of non negative IWCS.

2 Diagnosability with weights

Notation. For an alphabet A we note by A^* the set of finite sequences over A and by A^ω the set of infinite sequences. For $u \in A^*$ and $v \in A^*$ (resp. $v \in A^\omega$), u is a *prefix* of v ($u \prec v$) if there exists $w \in A^*$ (resp. $w \in A^\omega$) such that $uw = v$. The length of $u = a_1 \cdots a_n \in A^*$ is $|u| = n$ while $|u| = \infty$ for $u \in A^\omega$. Let $u \in A^\omega \cup A^*$ and $n \leq |u|$. Then $u_{|n}$ denotes the prefix of length n of u. The empty word is denoted ε.

Given a non-negative integer K, we denote $[K] = \{0, \dots, K\}$ and $\mathbb{N}_\omega$ the set of non-negative integers (denoted $\mathbb{N}$) enlarged with the ordinal ω and $\mathbb{N}_{>0}$ the set of positive integers. Similarly $\mathbb{Z}_\omega$ is the set of integers (denoted $\mathbb{Z}$) enlarged with the ordinal ω.

We study *weighted controllable transition systems* whose states are partitioned into correct and faulty ones with no transition from a faulty to a correct state. Every state has at least one outgoing transition. Transitions are labelled with actions, which can be controllable or not. Transitions between faulty states have an integer weight.

▶ **Definition 1.** *A* weighted controllable transition system *(WCS) $\mathcal{S} = \langle Q, \Sigma, T, W \rangle$ includes:*

- *a finite set of* states *$Q = C \uplus F$, where C are the correct states and F the faulty states;*
- *a finite set of* actions *$\Sigma = \Sigma_c \uplus \Sigma_u$ where Σ_c is the set of controllable actions and Σ_u is the set of uncontrollable actions;*
- *a set of* transitions *$T \subseteq Q \times \Sigma \times Q$ such that*
  - *for all $\langle q, a, q' \rangle \in T$, $q \in F \Rightarrow q' \in F$;*
  - *and for all $q \in Q$ there exist $a \in \Sigma$ and $q' \in Q$ such that $\langle q, a, q' \rangle \in T$;*
- *a* weight function *$W \colon T_F \to \mathbb{Z}$ with $T_F := T \cap (F \times \Sigma \times F)$. We call $\mathcal{S}$* non-negative *(resp.* positive*) if the co-domain of W is $\mathbb{N}$ (resp. $\mathbb{N}_{>0}$).*

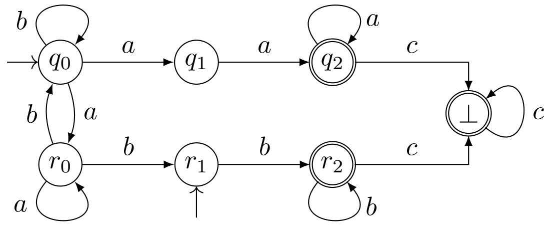

Figure 1 Example of a weight-observable IWCS with $W(a) = W(b) = 1$ and $W(c) = 0$.

When for all $a \in \Sigma$ and all pairs $\langle q, a, q' \rangle, \langle q_1, a, q'_1 \rangle \in T_F$ we have $W(\langle q, a, q' \rangle) = W(\langle q_1, a, q'_1 \rangle)$, we say that the WCS is weight-observable. *We then consider the partial function $W : \Sigma \to \mathbb{Z}$ defined by $W(a) = W(\langle q, a, q' \rangle)$ for actions a occurring within T_F.*

▶ **Example 2.** As a running example, we use the weight-observable non-negative WCS in Figure 1, where all actions are controllable, $W(a) = W(b) = 1$ and $W(c) = 0$. Generally, in weight-observable IWCS, we omit weights on transitions and specify W separately. In non-weight-observable IWCS, all transitions not explicitly equipped with a weight are assumed to have weight zero. The faulty states are depicted by a double circle.

Notation. We note $q \xrightarrow{a} q'$ for $\langle q, a, q' \rangle \in T$. Let $Q_1 \subseteq Q$ and $a \in \Sigma$, we define $\delta(Q_1, a) = \{ q' \in Q \mid \exists q \in Q_1 : q \xrightarrow{a} q' \}$ denoted by $Q_1 \xRightarrow{a} \delta(Q_1, a)$. Similarly $\delta^{-1}(q, a) = \{ q' \in Q \mid q' \xrightarrow{a} q \}$ and $\delta^{-1}(Q_1, a) = \bigcup_{q \in Q_1} \delta^{-1}(q, a)$. Let $w \in \Sigma^*$, then $\delta(Q_1, w)$ is inductively defined by $\delta(Q_1, \varepsilon) = Q_1$ and for all $a \in \Sigma$, $\delta(Q_1, wa) = \delta(\delta(Q_1, w), a)$. We assume that there exists $t \in T_F$ such that $W(t) \neq 0$ and define $0 < W_{\max} = \max\{ |W(t)| \mid t \in T_F \}$.

An item of $2^Q \setminus \emptyset$ is said to be a *belief*. A belief *Bel* is said to be *safe* if $Bel \subseteq C$ and *conclusive* if $Bel \subseteq F$. A belief *Bel* is said to be *ambiguous* if $Bel \cap C \neq \emptyset$ and $Bel \cap F \neq \emptyset$.

In an initialized WCS, the initial state may be unknown but is some state of a safe belief:

▶ **Definition 3.** *An initialized weighted controllable transition system (IWCS) is a pair $\langle \mathcal{S}, I \rangle$ where $\mathcal{S}$ is a WCS and I is a safe belief of $\mathcal{S}$.*

When representing an IWCS graphically, we shall indicate the states of I by incoming arrows. Figure 1 represents an IWCS with $I = \{q_0, r_1\}$. The next definition introduces runs and their traces, and partitions traces into surely faulty, surely correct and ambiguous ones.

▶ **Definition 4.** *Let $\langle \mathcal{S}, I \rangle$ be an IWCS. Then $\rho = q_0(a_n, q_n)_{0<n\leq K}$ is a* finite run *of $\langle \mathcal{S}, I \rangle$ if $q_0 \in I$ and for all $n < K$, $q_n \xrightarrow{a_{n+1}} q_{n+1}$. Its* trace *$tr(\rho)$ is the word $a_1 \ldots a_K$.*

- *ρ is* correct *if $q_K \in C$, otherwise ρ is* faulty*;*
- *$tr(\rho)$ is* surely correct *if $\delta(I, tr(\rho))$ is safe;*
- *$tr(\rho)$ is* surely faulty *if $\delta(I, tr(\rho))$ is conclusive;*
- *$tr(\rho)$ is* ambiguous *if $\delta(I, tr(\rho))$ is ambiguous.*

For instance, in Figure 1, the trace a is surely correct, aa is ambiguous, while aac surely faulty (and minimal at that). We also extend the notions to infinite runs and their traces.

▶ **Definition 5.** *Let $\langle \mathcal{S}, I \rangle$ be an IWCS. $\rho = q_0(a_n, q_n)_{n>0} \in (Q \times \Sigma)^\omega$ is a* run *of $\langle \mathcal{S}, I \rangle$ if $q_0 \in I$ and for all $K \in \mathbb{N}$, $\rho_K = q_0(a_n, q_n)_{0<n\leq K}$ is a finite run of $\langle \mathcal{S}, I \rangle$. Its trace is $tr(\rho) = (a_n)_{n>0} \in \Sigma^\omega$.*

- *ρ is* correct *if for all $K \in \mathbb{N}$, ρ_K is correct, otherwise ρ is* faulty*;*
- *$tr(\rho)$ is* surely correct *if for all $K \in \mathbb{N}$ there exists $K' > K$ with $tr(\rho_{K'})$ surely correct;*
- *$tr(\rho)$ is* surely faulty *if there exists K such that $tr(\rho_K)$ is surely faulty (which implies that for all $k \geq K$, $tr(\rho_k)$ is surely faulty);*
- *$tr(\rho)$ is* ambiguous *if it is neither surely correct nor surely faulty.*

In Figure 1, the trace $(aabb)^\omega$ is surely correct, any trace containing c is surely faulty, and b^ω is ambiguous. When a finite or infinite trace $w = (a_n)_{n<\alpha}$ with $\alpha \in \mathbb{N}_\omega$ is surely faulty, one denotes $\overline{w} = (a_n)_{n \leq K}$ the smallest prefix $a_1 \ldots a_K$ which is surely faulty. The weight of a faulty run whose trace is surely faulty is the sum of the weights of transitions that occur between faulty states until the corresponding trace is surely faulty.

▶ **Definition 6.** *Let $\langle \mathcal{S}, I \rangle$ be an IWCS and $\rho = q_0(a_n, q_n)_{0<n<\alpha}$ with $\alpha \in \mathbb{N}_\omega$ be a faulty run of $\langle \mathcal{S}, I \rangle$. Let $m = \min\{\, i \mid q_i \in F \,\}$ and $n = |\overline{tr(\rho)}|$. Then the* weight *of ρ is defined as $\kappa(\rho) = \sum_{m<i\leq n} W(\langle q_{i-1}, a_i, q_i \rangle)$.*

Observe that if the fault is immediately detected then $n = m$ and $\kappa(\rho) = 0$. The cost (resp. reward) of a surely faulty trace is the maximum (resp. minimum) weight of a faulty run that has triggered this trace.

▶ **Definition 7.** *Let $\langle \mathcal{S}, I \rangle$ be an IWCS and w be a surely faulty trace of $\langle \mathcal{S}, I \rangle$. Then:*

- *$Co(w)$, the* cost *of w, is defined by $Co(w) = \max\{\, \kappa(\rho) \mid \rho$ is a run of $\langle \mathcal{S}, I \rangle$ and $tr(\rho) = w \,\}$;*
- *$Re(w)$, the* reward *of w, is defined by $Re(w) = \min\{\, \kappa(\rho) \mid \rho$ is a run of $\langle \mathcal{S}, I \rangle$ and $tr(\rho) = w \,\}$.*

For instance, consider the trace $w = a^4 c^\omega$ in Example 2. We have $\overline{w} = a^4 c$ and $Co(w) = 2$ because the only run with this trace enters a faulty state after the second a.

The next definitions establish the notions of *controller* and *active diagnoser*. A controller observes a finite trace and deduces a belief, i.e., the set of possible current states. After each observed action, it selects a subset of actions that are allowed in the next step. This subset must be *admissible* w.r.t. the current belief, i.e. include the uncontrollable actions and allow at least one action enabled in each state of the belief. A controller is called *active diagnoser* if for any infinite faulty (resp. correct) run conforming to its control, the associated trace is surely faulty (resp. surely correct).

▶ **Definition 8.** *Let $\langle \mathcal{S}, I \rangle$ be an IWCS and Bel be a belief. A set of actions $\Sigma' \subseteq \Sigma$ is* admissible *for Bel if $\Sigma_u \subseteq \Sigma'$ (an uncontrollable action cannot be prevented) and for all $q \in Bel$, there exists $a \in \Sigma'$ such that $\delta(q, a) \neq \emptyset$.*

Notation. Let $\Sigma' \subseteq \Sigma$ and Bel be a belief. Then the set of enabled actions from Bel under control Σ', $En(Bel, \Sigma') \subseteq \Sigma'$, is defined by $\{a \in \Sigma' \mid \delta(Bel, a) \neq \emptyset\}$.

▶ **Definition 9.** *Let $\langle \mathcal{S}, I \rangle$ be an IWCS and σ a mapping from Σ^* to 2^Σ. Define $\mathcal{L}^f_\sigma \subseteq \Sigma^*$ as the set of words $w = a_1 \ldots a_n$ such that $\delta(I, w) \neq \emptyset$ and $a_{m+1} \in \sigma(w_{|m})$ for all $0 \leq m < n$. Then σ is a* controller *if for all $w \in \mathcal{L}^f_\sigma$, $\sigma(w)$ is admissible for $\delta(I, w)$.*

Thus, L^f_σ are the finite traces allowed by σ. We note $\mathcal{L}_\sigma \subseteq \Sigma^\omega$ as the set of all infinite words w such that all finite prefixes of w are in $\mathcal{L}^f_\sigma$. If σ is a controller, then $\mathcal{L}^f_\sigma$ is the set of finite prefixes of words of $\mathcal{L}_\sigma$, since σ cannot ever prevent a run from progressing.

▶ **Definition 10.** *Let $\langle \mathcal{S}, I \rangle$ be an IWCS and σ be a controller. Then σ is an* active diagnoser *if for all $w \in \mathcal{L}_\sigma$, w is either surely correct or surely faulty. We call $\langle \mathcal{S}, I \rangle$* actively diagnosable *if an active diagnoser exists for it.*

Example 2 (continued). For the system of Figure 1, we describe an active diagnoser achieving an arbitrary reward B (*i.e.* such that every surely faulty trace has reward at least B). Initially, it only allows action b, which leads to $\{q_0, r_2^0\}$ (the superscript 0 means that the current reward associated with this state is zero). Then again it only allows B times action b which leads to $\{q_0, r_2^B\}$. Afterwards it allows $\{a, c\}$. If a is observed, this leads to $\{q_1, r_0\}$, a situation symmetrical to the initial one where a similar strategy can be played with action a instead of action b and vice versa. If c is observed, this leads to $\{\perp^B\}$ where a fault is diagnosed with reward B.

Decision problems

In the sequel, we address three relevant problems for active diagnosis with weights.

1. **The cost problem.** Let $\langle \mathcal{S}, I \rangle$ be an IWCS and B be an integer. Does there exist an active diagnoser σ such that for all surely faulty $w \in \mathcal{L}_\sigma$, $Co(w) \leq B$?
2. **The reward problem.** Let $\langle \mathcal{S}, I \rangle$ be an IWCS and B be an integer. Does there exist an active diagnoser σ such that for all surely faulty $w \in \mathcal{L}_\sigma$, $Re(w) \geq B$?
3. **The unbounded reward problem.** Let $\langle \mathcal{S}, I \rangle$ be an IWCS. For all integers B does there exist an active diagnoser σ such that for all surely faulty $w \in \mathcal{L}_\sigma$, $Re(w) \geq B$?

In the following, we will study these problems for various restrictions of IWCS:

▶ **Definition 11.** *An IWCS $\langle \mathcal{S}, I \rangle$ has a* single fault *(*SF-IWCS*) if there is $\iota \in \Sigma$ such that:*
- *for all $\langle q, a, q' \rangle \in T \cap (C \times \Sigma \times F)$, $a = \iota$;*
- *for all $\langle q, a, q' \rangle \in T_F$, $a \neq \iota$.*

A SF-IWCS $\langle \mathcal{S}, I \rangle$ has a single initial fault *(*SIF-IWCS*) if there exists $q_0 \in C$ such that:*
- *$I = \{q_0\}$ and for all $\langle q_0, a, q \rangle \in T$, $a = \iota$;*
- *for all $\langle q, a, q' \rangle \in T$, $q' \neq q_0$ and $a = \iota \Rightarrow q = q_0$.*

Notations. In other words, in a SF-IWCS, the only transitions that go to a faulty state from a correct one are labelled by ι, and moreover ι does not label faulty transitions. This implies that when ι is observed the system was correct before the observation and may have become faulty at that moment. Moreover, in a SIF-IWCS a fault can happen during the very first action ι but not afterwards. So we introduce its initial relevant belief $I^+ = \delta(q_0, \iota)$ and may consider that the system starts in I^+.

Observations. For the purposes of our three problems, the control of an active diagnoser is irrelevant once it reaches a conclusive belief. Therefore we assume w.l.o.g. that afterwards the admissible set for control is always Σ. In SIF-IWCS, the same assumption can be made for safe beliefs. Observe that when the weights lie in $\mathbb{Z}$, the cost and the reward problems are mutually reducible by simultaneously multiplying all the weights and the threshold by -1. This reduction is no longer possible as soon as these weights are restricted to $\mathbb{N}$. Since the existence of an active diagnoser without weights is already EXPTIME-hard (see [5] for the non deterministic framework and [1] for the probabilistic framework), this lower bound applies to all our decidability results.

3 Undecidability of all problems for SIF-IWCS

In this section, we show that the cost, reward, and unbounded reward problems are all undecidable in general, in particular when the weights are from $\mathbb{Z}$ and the system is not weight-observable, even for single initial faults. As said above, the cost and reward problems are equivalent with integer weights. Thus, we only need to treat the case of rewards.

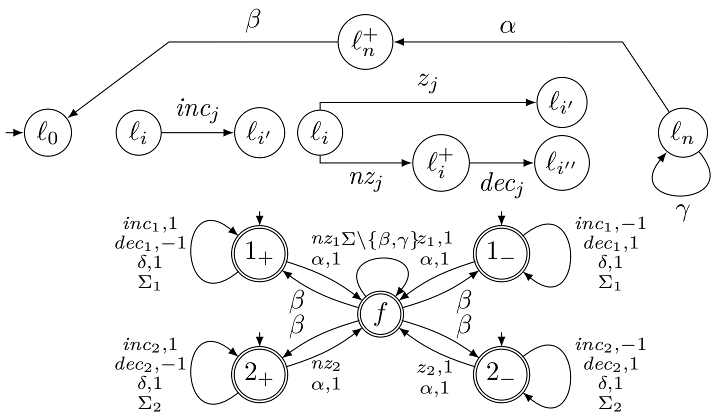

Figure 2 The simulation of the counter machine, where $\Sigma_j = \Sigma \setminus \{inc_j, dec_j, \alpha, \beta, \gamma, \delta\}$.

Notation. When depicting a SIF-IWCS in a figure, we omit the initial state and q_0 its outgoing transitions and represent the states of I^+ as initial ones with an input edge. The label of a transition is a pair ⟨action,weight⟩ with the weight omitted when null.

▶ **Theorem 12.** *The cost, reward, and unbounded reward problems for SIF-IWCS are undecidable.*

Sketch of proof. We simultaneously reduce the halting problem of a two-counter machine $\mathcal{C}$ to the reward and unbounded reward problems of an IWCS $\langle \mathcal{S}, I\rangle$. More precisely, the reduction produces a SIF-IWCS such that if $\mathcal{C}$ halts then for all B there is an active diagnoser that achieves a reward B, otherwise there is no active diagnoser that achieves even a reward of 1. The two counters of $\mathcal{C}$ are c_1 and c_2, and its set of locations is $\{\ell_0, \dots, \ell_n\}$, where ℓ_0 is the initial and ℓ_n the final location. Each location ℓ_i has one of two operations on some counter c_j: increment (inc_j; **goto** $\ell_{i'}$) or test-and-decrement (**if** $c_j = 0$ **then goto** $\ell_{i'}$ **else** dec_j; **goto** $\ell_{i''}$). W.l.o.g. we assume that on termination the two counters are reset to zero and that there is at least one test-and-decrement instruction.

The simulation of $\mathcal{C}$ by the SIF-IWCS $\langle \mathcal{S}, I\rangle$ is depicted in Figure 2, where $I = \{\ell_0, 1_-, 1_+, 2_-, 2_+\}$. The set of states is defined by $C = \{q_0, \ell_0, \dots, \ell_n\} \cup \{\ell_i^+ \mid i = n \text{ or instruction } i \text{ is a test}\}$ and $F = \{1_-, 1_+, 2_-, 2_+, f\}$. The set of actions is $\Sigma = \{inc_j, dec_j, z_j, nz_j \mid i \in \{1,2\}\} \cup \{\alpha, \beta, \gamma, \delta\}$. All the actions are controllable.

The upper part of Figure 2 represents the correct runs of the IWCS. The transitions from ℓ_n via ℓ_n^+ to ℓ_0 allow to repeat the simulation. The simulation of $\mathcal{C}$ is weak since it does not take into account the values of the counters for the test instructions. The lower part of Figure 2 represents faulty runs of the IWCS. Intuitively, the faulty runs in j_+ and j_- both represent the counter c_j, with the first one having weight c_j (on the first iteration of $\mathcal{C}$ at least) and the second having weight $-c_j$. To do so, when the current state is j_+ (resp. j_-) the loop transition labeled by inc_j increments (resp. decrements) the reward of the run while the loop transition labeled by dec_j decrements (resp. increments) the reward of the run. When going through a test, the controller decides whether to allow z_j (which claims the counter c_j is equal to 0), or nz_j. In the first case, the faulty run in j_- should have weight 0, and thus adding 1 while going to f ensures a weight of 1. If the counter was not equal to 0 however, this run has a negative weight, and thus allowing z_j means a run with weight at most 0 reached f. In the case of nz_j, we similarly allow the run in j_+ to reach f (without any increase this time), and this run has weight at least 1 iff the counter was not 0. In other words, the weight of the faulty runs in f is at least 1 iff the active diagnoser correctly selected its actions z_j or nz_j depending on whether the counter c_j was equal or not to 0. Ignoring the ability to restart the system from l_n which is useful for the unbounded cost problem, the active diagnoser can only detect the fault once ℓ_n is reached, and the faulty runs have weight at least 1 iff the runs ending in f have weight at least 1, and thus iff the two-counter machine was correctly emulated. ◀

4 Decidability of the problems for restricted IWCS

In Section 3, we established that our problems of interest are undecidable in general. We now identify two restrictions under which the problems do remain decidable, establish upper bounds for their complexity, and provide game-theoretic methods for synthesizing active diagnosers that achieve given costs or rewards. In Section 4.1, we allow arbitrary weights but impose restrictions on observability and structure (notably, single fault). In Section 4.2, we impose a restriction on weights (non negative) but none on observability or structure. Section 4.3 provides lower bounds on the memory size of active diagnosers achieving given costs or rewards.

To begin with, the following lemma about the existence of active diagnosers was proved in the framework of controllable systems *without* weights. Here, we establish a lower bound on its reward.

▶ **Lemma 13.** *Let $\mathcal{S}$ be a WCS and I be an initial, not necessarily safe belief. Then one can decide in* EXPTIME *whether there exists an active diagnoser for $\langle \mathcal{S}, I\rangle$. Furthermore, if this active diagnoser exists, the reward of this active diagnoser is at least $-2^{|Q|}W_{\max}$, and its cost is at most $2^{|Q|}W_{\max}$.*

4.1 Weight-observable Single Fault IWCS

We first study the case where weights can be drawn from $\mathbb{Z}$, but with two restrictions: (i) the system must be weight-observable, and (ii) the system has a single type of fault. In this case, all three problems are solvable in exponential time (Theorems 21 and 23).

Notation and observation. Let $w = w_0\iota w_1$ with $w_1 \in (\Sigma \setminus \{\iota\})^*$ be a finite ambiguous trace of a weight-observable SF-IWCS. Let us define $\kappa(w) = W(w_1)$. Then $\kappa(w)$ is the common value of $\kappa(\rho)$ for all faulty runs ρ such that $tr(\rho) = w$.

▶ **Definition 14** (The computation tree). *Let $\langle \mathcal{S}, I\rangle$ be a weight-observable SF-IWCS and σ be a controller. Then $\Upsilon(\sigma)$, the (possibly infinite)* computation tree *of σ, is inductively defined as follows.*

- *Its edges are labeled by characters in Σ such that two edges outgoing from a vertex have different labels. Thus every vertex v in the tree has for identifier the label of the path from the root $r = \varepsilon$ to v;*
- *Its root is labelled by $\lambda(r) = \langle \omega, I\rangle \in \mathbb{Z}_\omega \times 2^Q$;*
- *Let v be a vertex labelled by $\lambda(v) = \langle k, Q'\rangle$ with $Q' \cap C \neq \emptyset$. Let $a \in En(Q', \sigma(v))$ and $H = \delta(Q', a)$. Then there is an edge $v \overset{a}{\Rightarrow} va$ where:*
  1. *if $H \subseteq C$ then $\lambda(va) = \langle \omega, H\rangle$;*
  2. *else if $a = \iota$ then $\lambda(va) = \langle 0, H\rangle$;*
  3. *else $\lambda(va) = \langle k + W(a), H\rangle$.*
- *Let v be a vertex such that $\lambda(v) \in \mathbb{Z} \times 2^F$. Then v is a leaf.*

Notation. Let $\ell = \langle k, Q' \rangle \in \mathbb{Z}_\omega \times 2^Q$. Then $\ell^\downarrow$ denotes the belief Q'.

Observation. If σ is an active diagnoser, then its cost (resp. reward) is equal to $\sup(k \mid \exists v \text{ leaf of } \Upsilon(\sigma)\ \lambda(v) \in \{k\} \times 2^F)$ (resp. $\inf(k \mid \exists v \text{ leaf of } \Upsilon(\sigma)\ \lambda(v) \in \{k\} \times 2^F)$) with the convention $\sup(\emptyset) = -\infty$ (resp. $\inf(\emptyset) = \infty$).

▶ **Definition 15** (Component of the computation tree)**.** *Let $\langle \mathcal{S}, I \rangle$ be a weight-observable SF-IWCS, σ be a controller and v be a vertex of $\Upsilon(\sigma)$ such that $\lambda(v)^\downarrow \in 2^C$. Then the* v-component *of $\Upsilon(\sigma)$, denoted $\Upsilon(\sigma)[v]$, is the subtree of $\Upsilon(\sigma)$ rooted in v and inductively defined as follows. Let w be a vertex of $\Upsilon(\sigma)[v]$.*

- *If either $w = v$ or $\lambda(w)^\downarrow \notin 2^C$ then for all $a \in \Sigma$ there is an edge $w \stackrel{a}{\Longrightarrow} wa$ in $\Upsilon(\sigma)[v]$ iff there is an edge $w \stackrel{a}{\Longrightarrow} wa$ in $\Upsilon(\sigma)$;*
- *Else w is a leaf.*

Observation. $\Upsilon(\sigma)$ can be viewed as the "concatenation" of all $\Upsilon(\sigma)[v]$.

Active diagnosers produce only finite components of their computation tree as established in the following Lemma.

▶ **Lemma 16.** *Let $\langle \mathcal{S}, I \rangle$ be a weight-observable SF-IWCS and σ be a controller. Then σ is an active diagnoser if and only if all the components of its computation tree are finite.*

A repetitive active diagnoser applies the same control on components whose roots have the same label; the active diagnoser's choice depends only on the current belief. It is efficient if the weight of every subrun from a belief to itself without the occurrence of a fault (witnessed by the occurrence of ι) is positive.

▶ **Definition 17.** *Let $\langle \mathcal{S}, I \rangle$ be a weight-observable SF-IWCS and σ be an active diagnoser. Then σ is* repetitive *if for every two roots v and v' of components of $\Upsilon(\sigma)$ with $\lambda(v) = \lambda(v')$, $\Upsilon(\sigma)[v]$ and $\Upsilon(\sigma)[v']$ are isomorphic via the mapping $f(vz) = v'z$.*

Furthermore σ is efficient *if there does not exist in any component $\Upsilon(\sigma)[v_{Q'}]$, vertices v and vv' with $v' \in (\Sigma \setminus \{\iota\})^+$ fulfilling $W(v') \leq 0$ and $\lambda(v)^\downarrow = \lambda(vv')^\downarrow$.*

The next proposition shows that w.r.t. the three decision problems one can restrict attention to the set of repetitive efficient active diagnosers.

▶ **Proposition 18.** *Let $(\mathcal{S}, I)$ be a weight-observable SF-IWCS and σ be a active diagnoser such that for all surely faulty $w \in \mathcal{L}_\sigma$, $Re(w) \geq B$. Then there exists a repetitive efficient active diagnoser σ' such that for all surely faulty $w \in \mathcal{L}_{\sigma'}$, $Re(w) \geq B$.*

Due to Lemma 13, when an IWCS is actively diagnosable, there exists an active diagnoser with a reward at least $B' := -2^{|Q|} W_{\max}$. So when the threshold is smaller than B', the reward problem boils down simply to the existence of any active diagnoser, which is solvable in exponential time. Given $B \geq B'$, we construct a sequential Büchi game with players D (diagnoser) and E (environment) such that the existence of a winning strategy for D is equivalent to the existence of active diagnoser achieving reward B. A location owned by D is either the absorbing state $\top$ corresponding to a win, or the absorbing state $\bot$ corresponding to a loss, or consists of a pair whose first item is the reward of the current trace and the second item is the set of possible states. The actions of D in this location are exactly the admissible controls w.r.t. the possible states. The locations of E are the locations of D (except $\top$ and $\bot$) paired with an admissible control, and E can choose among the enabled and admissible actions, leading to (1) a loss if this results in a surely faulty trace with reward less than B, or an ambiguous trace with reward less than B', (2) a win if all possible states are faulty and the reward is at least B, or (3) otherwise a location of D where its two items are updated according to the action. The winning locations for D are $\top$ and the locations with no faulty states (as such, the Büchi condition ensures that an infinite trace through these locations is surely correct).

▶ **Definition 19.** *Let $\langle \mathcal{S}, I\rangle$ be a weight-observable SF-IWCS and $-2^{|Q|}W_{\max} \leq B$. Then the sequential Büchi* reward game *$\mathcal{G}_r(\mathcal{S}, I, B) = \langle L, \ell_0, \Delta, Acc\rangle$ between D(iagnoser) and E(nvironment) is inductively defined by:*

- *$L = L_E \uplus L_D$ is the set of locations, where:*
  - *$L_D \subseteq \big(\{\infty\} \cup [-2^{|Q|}W_{\max}, B + 2^{|Q|}W_{\max}]\big) \times 2^Q \ \cup\ \{\top, \bot\}$ is the set of states controlled by D;*
  - *$L_E \subseteq \big(\{\infty\} \cup [-2^{|Q|}W_{\max}, B + 2^{|Q|}W_{\max}]\big) \times 2^Q \times 2^\Sigma$ those controlled by E;*
  - *$\ell_0 = \langle \infty, I\rangle$ is the initial location.*
- *Δ contains the transitions $\top \Rightarrow \top$ and $\bot \Rightarrow \bot$, making $\bot$ and $\top$ sink states;*
- *Let $\ell = \langle k, Q'\rangle \in L_D$. Then for all Σ' admissible for Q' there is a transition $\ell \overset{\Sigma'}{\Longrightarrow} \ell'$ with $\ell' = \langle k, Q', \Sigma'\rangle$ owned by E;*
- *Let $\ell = \langle k, Q', \Sigma'\rangle \in L_E$. Then for all $a \in En(Q', \Sigma')$ there is a transition $\ell \overset{a}{\Rightarrow} \ell'$ where ℓ' owned by D is defined as follows, where $Q'' = \delta(Q', a)$:*
  1. *$\ell' = \bot$ when:*
     a. *$Q'' \subseteq F$ and either (1) $a = \iota$ and $B > 0$ or (2) $a \neq \iota$ and $k + W(a) < B$;*
     b. *$Q'' \cap F \neq \emptyset$, $a \neq \iota$ and $k + W(a) < -2^{|Q|}W_{\max}$.*
  2. *$\ell' = \top$ when $Q'' \subseteq F$ and either (1) $a = \iota$ and $B \leq 0$ or (2) $a \neq \iota$ and $k + W(a) \geq B$;*
  3. *Otherwise $\ell' = \langle k', Q''\rangle$ where k' is equal to:*
     a. *∞ if $Q'' \subseteq 2^C$;*
     b. *0 if $Q'' \cap F \neq \emptyset$ and $a = \iota$;*
     c. *$\min(B + 2^{|Q|}W_{\max}, k + W(a))$ if $Q'' \cap F \neq \emptyset$ and $a \neq \iota$.*
- *Acc, the set of accepting states, is $\{\top\} \cup \{\langle \infty, Q'\rangle\}_{Q' \subseteq C}$.*

▶ **Proposition 20.** *Let $\langle \mathcal{S}, I\rangle$ be a weight-observable SF-IWCS and $-2^{|Q|}W_{\max} \leq B$. There exists an active diagnoser σ such that for all surely faulty $w \in \mathcal{L}_\sigma$, $Re(w) \geq B$ if and only if D has a winning strategy in $\mathcal{G}_r(\mathcal{S}, I, B)$.*

Given a sequential Büchi game $\mathcal{G}$ and a player P, deciding whether P has a winning strategy for $\mathcal{G}$ can be done in polynomial time (see for instance [9]). Thus since $\mathcal{G}_r(\mathcal{S}, I, B)$ has exponential size w.r.t. the size of the reward problem, one gets:

▶ **Theorem 21.** *The reward problem for weight-observable SF-IWCS can be decided in exponential time.*

▶ Remark. The size of $\mathcal{G}_r(\mathcal{S}, I, B)$ is still a single exponential w.r.t the size of $\mathcal{S}$ even if B is exponential w.r.t. the size of $\mathcal{S}$. Furthermore, since a winning strategy in $\mathcal{G}_r(\mathcal{S}, I, B)$ can be chosen memoryless, the memory size of the active diagnoser is the number of states owned by D associated with an admissible control, which is in $\Theta((2^{|Q|} \cdot W_{\max} + B)|\Sigma|)$.

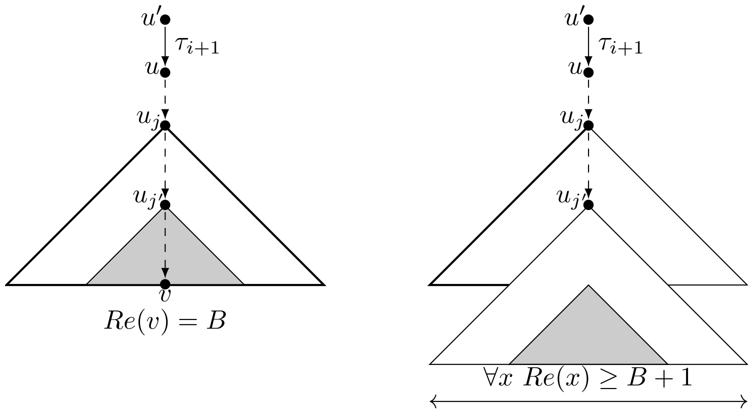

Figure 3 Expanding an active diagnoser.

The next proposition states that given a weight-observable SF-IWCS, if there exists an active diagnoser achieving a reward beyond a certain threshold, where that threshold depends only on the IWCS, then for any reward B there exists an active diagnoser achieving it.

▶ **Proposition 22.** *Let $(\mathcal{S}, I)$ be a weight-observable SF-IWCS and $B \geq 2^{|Q|}W_{\max}$. If there exists an active diagnoser σ such that for all surely faulty $w \in \mathcal{L}_\sigma$, $Re(w) \geq B$ then there exists an active diagnoser σ' such that for all surely faulty $w \in \mathcal{L}_{\sigma'}$, $Re(w) \geq B + 1$.*

Sketch of proof. Due to Proposition 18 w.l.o.g. assume σ is repetitive and efficient. We build σ' by modifying each component of $\Upsilon(\sigma)$ independently.

Fix a component $\Upsilon(\sigma)[v_{Q'}]$. An edge $v_1 \xRightarrow{\iota} v_2$ is deemed *critical.* Let $Cr = \{\tau_1, \ldots, \tau_k\}$ be its critical edges, ordered so that if τ_i is a descendant of τ_j, then $i < j$. For a faulty leaf v, let $\tau(v)$ denote the last critical edge on the path from $v_{Q'}$ to v.

We inductively build $\sigma = \sigma_0, \ldots, \sigma_k$, modifying only the $v_{Q'}$-component, so that σ_i satisfies: for every faulty leaf v, (a) if $\tau(v) \in Cr_\ell$ then $Re(v) \geq B + 1$, and (b) if $\tau(v) \in Cr_u = \{\tau_{i+1}, \ldots, \tau_k\}$ then $Re(v) \geq B$.

Induction step. Suppose σ_i satisfies the above and there exists a faulty leaf v with $\tau(v) = \tau_{i+1}$ and $Re(v) = B$. Let u be the target of τ_{i+1}. Since no ι occurs on the path from u to v, among the $2^{|Q|} + 1$ prefixes $u_0, \ldots, u_{2^{|Q|}}$ of v beyond u (with $u_0 = u$ and $\lambda(u_j) = \langle k_j, Q_j \rangle$, $k_j \in [jW_{\max}, (j{+}1)W_{\max}[$), there exist $j < j'$ with $Q_j = Q_{j'}$ and $k_j < k_{j'}$ (left part of Figure 3). Define $\overline{\sigma}$ as σ_i except that every $u_{j'}w$ follows the control of u_jw: this inserts a copy of the sub-tree between u_j and $u_{j'}$ (right part of Figure 3), strictly increasing Re of the affected leaves. Since the number of faulty leaves with reward exactly B decreases, the process terminates and yields σ_{i+1}.

After k steps, $Cr_u = \emptyset$ and all faulty leaves of $\Upsilon(\sigma_k)[v_{Q'}]$ satisfy $Re(v) \geq B+1$. Applying this to every component produces σ' with $Re(w) \geq B + 1$ for all minimal surely faulty $w \in \mathcal{L}_{\sigma'}$. ◀

Due to Proposition 22, the unbounded reward problem for SF-IWCS is equivalent to solving the reward problem for a threshold of $2^{|Q|}W_{\max}$. Using Theorem 21 and the following remark this can be done in EXPTIME.

▶ **Theorem 23.** *The unbounded reward problem for weight-observable SF-IWCS can be decided in exponential time.*

4.2 Non negative IWCS

We now turn to the case of non negative IWCS. Similar to Section 4.1, we construct a Büchi game such that the existence of active diagnoser whose cost (resp. reward) is at most (resp. at least) some threshold B is equivalent to the existence of a winning strategy for player D. Since solving a Büchi game can be done in polynomial time in the size of the Büchi game and this game is a single exponential of the cumulated size of $(\mathcal{S}, I)$ and $\log(B)$ one immediately gets:

▶ **Theorem 24.** *The cost problem can be decided in exponential time.*

▶ **Theorem 25.** *The reward problem can be decided in exponential time.*

In Proposition 22, we established how to solve the unbounded reward problem for SF-IWCS by extending an active diagnoser achieving bound B to an active diagnoser achieving bound $B + 1$. The unbounded reward problem for non negative weight-observable IWCS and for positive IWCS can be decided by establishing a similar result, though the proof is entirely different and far more complex.

▶ **Theorem 26.** *The unbounded reward problem for non-negative weight-observable IWCS and for positive IWCS can be decided in double exponential time.*

Sketch of proof. In Proposition 22, we selected a bound B large enough that in every path of a component of the computation tree of an active diagnoser ensuring B, the same set of states would be visited twice, with an increase of value in between. This allowed us to repeat the behaviour of the active diagnoser after the second occurrence, which will ensure a value of at least $B + 1$, proving that achieving B means one can ensure any value.

In this proof, we try to follow the same idea. However, having read a word of positive weight does not guarantee an increase of the values of the runs: while reading this word, a new faulty run may have been created which weight naturally starts at 0 and thus might be lower than what the system started with. Hence, our first step is to interpret our IWCS as a vector of values, where each coordinate correspond to a state of the system, and the value in this coordinate correspond to the lowest weight a faulty run ending in this state could have. We establish a (doubly exponential) bound on the length of antichains in this representation, which indicates how long a branch of the computation tree must be to have a pair of configurations we can repeat to increase the weights of runs.

This is however not sufficient to increase the reward on this branch, as due to the creation of new faults, we might need multiple repetitions of the strategy of the active diagnoser throughout a single branch. However, a careful analysis of the impact of the repetition allows us to bound how many times a repetition might be needed within a branch, finally producing the active diagnoser increasing the value of the bound. ◀

4.3 Size complexity of the active diagnosers

In this section, we establish lower bounds for the size of controllers required to perform active diagnosis. We first introduce the following definition:

▶ **Definition 27** (state-based controller)**.** *Let $(\mathcal{S}, I)$ be an IWCS and σ be a controller. Then σ is* state-based *if there exists a right-congruence $\sim$ over $\mathcal{L}_\sigma^f$ with a finite quotient $L_\sigma^f/\sim$ such that for all $w \sim w'$, $\sigma(w) = \sigma(w')$. The size of σ is said to be $|L_\sigma^f/\sim|$.*

A state-based controller can be implemented as a finite-state machine that reads an observation and decides its control based on its current state. We will study the necessary number of states of such machines.

Costs

The cost problem is closely related to the notion of *delay* studied in [5]. Let $\langle \mathcal{S}, I \rangle$ be an IWCS such that all weights are 1 and σ an active diagnoser. In this case, the *delay* of σ is the maximal cost $Co(w)$ of a surely faulty trace $w \in \mathcal{L}_\sigma$.

▶ **Proposition 28** ([5], Theorem 4). *There is a family $(\mathcal{A}_n)_{n\geq 1}$ of actively-diagnosable IWCS such that (i) $\mathcal{A}_n$ is of size $\mathcal{O}(n)$, (ii) there exist active diagnosers with delay n, and (iii) any such diagnoser has at least size $n!$.*

Applying Proposition 28 to IWCS where the weight function is constant and equal to 1 (thus, the cost of a run is exactly the delay between the fault occurrence and its detection), we obtain that there exists a family of IWCS such that any active diagnoser with $Co(\sigma) \leq n$ has at least size $n!$. Moreover, there exists a family of systems with exponential intrinsic delay:

▶ **Proposition 29** ([5], Theorem 5). *There is a family $(\mathcal{A}'_n)_{n\geq 1}$ of SIF-IWCS where (i) $\mathcal{A}'_n$ has $\mathcal{O}(n)$ states; (ii) $\mathcal{A}'_n$ is actively diagnosable, and (iii) every active diagnoser has delay at least 2^n.*

From Proposition 29, it follows immediately that there exists a family of SIF-IWCS such that any active diagnoser satisfies $Co(\sigma) \geq 2^n \cdot W_{\max}$. Moreover, Lemma 13 implies that this lower bound is tight.

Rewards

We now study active diagnosers that attempt to obtain a reward B for an IWCS with n states. It turns out that lower bounds for their size depend on the characteristics of the IWCS. These are summarized in the following table, where W_{max} and $|\Sigma|$ are assumed constant.

weight observability	weight-observable	general	
weight domain	general	non negative	
structure	single fault	general	
lower bound	$\Omega(2^n)$ for $B = 2^n$	$\Omega(n!)$ for $B = n$ even when weight-observable	$\Omega(B^n)$ even when SIF
upper bound	$\mathcal{O}(2^n \cdot B)$	$\mathcal{O}(B^n)$	

For the case of weight-observable SF-IWCS, observe that constructing an active diagnoser involves, among other things, building a determinized version of the IWCS in question. Thus, for a fixed B, a trivial lower bound for the size of an active diagnoser is $2^{\mathcal{O}(n)}$ (see Proposition 29 for an example achieving this lower bound). For weight-observable SF-IWCS, the upper bound comes from the size of the game constructed in the proof of Theorem 21.

We now discuss the family of weight-observable, non-negative IWCS shown in Figure 4. The lower bound established by this example is the same as for costs, but different techniques are required to prevent the environment from revealing a fault prematurely. We only sketch the intuition here.

We suppose that the diagnoser wants to achieve a reward of $n+1$ and set $\Sigma_c = \{a, b\} \cup \Delta$ and $\Sigma_u = \Gamma \cup \{e, f\}$, where $\Gamma = \{c_1, \ldots, c_n\}$ and $\Delta = \{d_1, \ldots, d_n\}$; moreover, $\Delta_i = \Delta \setminus \{d_i\}$. Only the controllable actions give a reward, which is 1. Intuitively, the diagnoser wants the first $2n+1$ actions to be $ac_{\pi(1)} \cdots ac_{\pi(n)}b$ since playing a or b at other moments would reveal a fault without the desired reward. In general, π is an endomorphism on $\{1, \ldots, n\}$, the most difficult case being a permutation, which results in belief $\{q, v^n_{\pi(1)}, \ldots, v^1_{\pi(n)}\}$. In

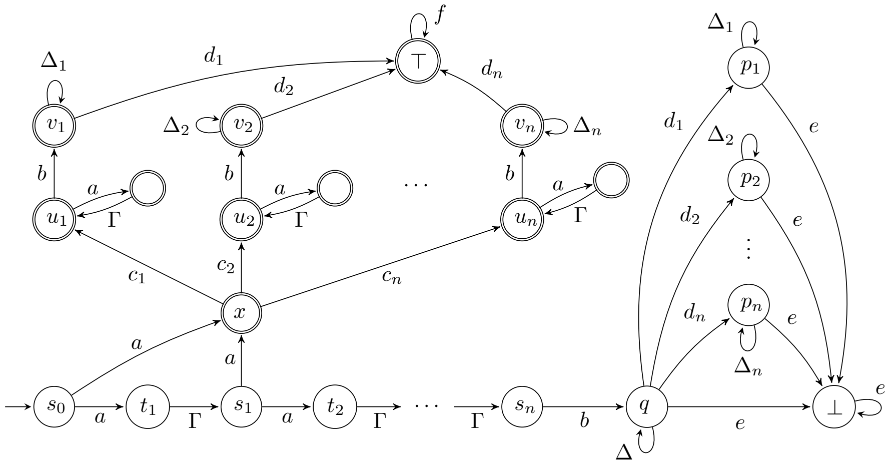

Figure 4 A weight-observable IWCS with $\mathcal{O}(n)$ states where the optimal reward requires an active diagnoser with at least $n!$ states.

order to guarantee its reward, the diagnoser must then allow $d_{\pi(1)}, \ldots, d_{\pi(n)}$, in that order. If the run is faulty, this will sooner or later send the system to $\top$ and reveal the fault by playing f, where the reward is always $n+1$. After this, the run can prove to be correct by playing e. The diagnoser must therefore, among other things, remember the order of the first n actions from Γ.

▶ **Proposition 30.** *There exists a family $(\mathcal{A}_n)_{n\geq 1}$ of weight-observable IWCS where $\mathcal{A}_n$ has $\mathcal{O}(n)$ states such that any state-based active diagnoser σ satisfying $Re(\sigma) \geq n+1$ is at least of size $n! = 2^{\mathcal{O}(n \log n)}$.*

In Section 4, a game of size $\mathcal{O}(B^n)$ is constructed to decide whether there exists an active diagnoser that obtains reward B, for a system with n faulty states. A winning strategy for this game immediately gives rise to a state-based active diagnoser of the same size.

We now turn to the rightmost column from the table above: When the system is not weight-observable, the lower bound of an active diagnoser for the reward problem can indeed be $\mathcal{O}(B^n)$ even for non-negative SIF-IWCS, where B is the reward the diagnoser strives to achieve.

▶ **Proposition 31.** *There exists a family $(\mathcal{B}_n)_{n\geq 1}$ of SIF-IWCS where $\mathcal{A}'_n$ has $\mathcal{O}(n)$ states such that any state-based active diagnoser σ satisfying $Re(\sigma) \geq B$, for any $B \geq 1$, is at least of size $(B+1)^n$.*

The upper bound for the latter two cases comes from the game constructed in the proof of Theorem 25.

5 Active diagnosability of pLTS with weights

In order to consider stochastic models, we need to extend the definitions of Section 2 to include probabilities.

▶ **Definition 32.** *A (resp. positive) probabilistic weighted controllable transition system (pWCS) $\mathcal{S}_{\mathcal{P}} = \langle Q, \Sigma, T, W, P\rangle$ is defined by:*

- *a WCS $\mathcal{S} = \langle Q, \Sigma, T, W\rangle$*
- *A probability function $P : T \to \mathbb{Q}_+$ with $\mathbb{Q}_+$ the set of positive rational number.*

An initialized weighted controllable probabilistic transition system (pIWCS) is a pair $(\mathcal{S}_\mathcal{P}, \mu)$ where $\mathcal{S}_\mathcal{P}$ is a pWCS and μ is a probability distribution which support is a safe belief of $\mathcal{S}_\mathcal{P}$.

Let $\mathcal{S}_\mathcal{P}$ be a pIWCS and σ be a controller. Probabilities of cylinders of runs can be defined thanks to the controller setting which actions are available:

$$\mathbb{P}_{\mathcal{S}_\mathcal{P},\sigma}(Cyl(q_0 a_0 q_1 \dots q_n)) = \mu(q_0) \times \prod_{i=0}^{n-1} \frac{P(q_i, a_i, q_{i+1})}{\sum_{q' \in Q, b \in \sigma(a_0 \dots a_{i-1})} P(q_i, b, q')} \ .$$

The probability measure is extended to measurable subsets by Caratheodory's extension theorem. To simplify notations, given a finite run ρ, we will sometimes abuse notation and write $\mathbb{P}(\rho)$ for $\mathbb{P}_{\mathcal{S}_\mathcal{P},\sigma}(Cyl(\rho))$ when the pIWCS and controller are clear from context. Similarly if R is a denumerable set of finite runs such that no run is a prefix of another one, $\mathbb{P}(R) = \sum_{\rho \in R} \mathbb{P}(\rho)$ (since all intersections of associated cylinders are empty).

In a pWCS, a controller is an almost surely active diagnoser, if with probability 1, $w \in \mathcal{L}_\sigma$ is either surely correct or surely faulty. In other words, $\mathbb{P}_{\mathcal{S}_\mathcal{P},\sigma}(\{w \in \mathcal{L}_\sigma \mid w$ is either surely correct or surely faulty$\}) = 1$. The decision problems for IWCS can straightforwardly be translated to pIWCS using the notion of almost surely active diagnoser.

We show that the frontier between decidability and undecidability is different for IWCS and pIWCS. More precisely we establish that contrary to the case of weight-observable SF-IWCS, the three decision problems for pIWCS are undecidable even when the corresponding IWCS is a weight-observable SIF-IWCS (i.e, a subclass of SF-IWCS).

▶ **Theorem 33.** *The stochastic cost problem, stochastic reward problem and stochastic unbounded reward problem for weight-observable pSIF-IWCS are undecidable.*

When restricted to non-negative weights, the stochastic cost and reward problems can be translated into an almost sure Büchi objective for a POMDP.

▶ **Theorem 34.** *The stochastic cost and reward problems for non negative pIWCS are decidable in exponential time.*

6 Conclusion and perspectives

We have associated weights with transitions of partially observed controllable systems in order to define costs and rewards for active diagnosers. This approach has led to three decision problems close to optimization problems like finding an optimal active diagnoser w.r.t. cost/rewards. We have first shown that these problems are undecidable even with restrictions as soon as the weights are integers and the system is not weight-observable. Thus we have proposed several incomparable restrictions all leading to decidability of some or all the problems with a complexity varying from EXPTIME to 2EXPTIME. Furthermore we established tight lower and upper bounds for the memory size of the active diagnosers below/above some threshold for the cost/reward requirement related to the size of the system and the value of the threshold.

Afterwards we have considered a probabilistic framework that allows to relax the diagnosis requirement to only hold almost surely. In this context, the problems become undecidable for integer weights even if the system is weight-observable. Then we have shown that two of the decision problems become decidable for weight-observable IWCS using natural integer weights.

Whilst we have established some frontier between decidability and undecidability, there are still problems to study whose status is unknown, for instance the unbounded reward problem for non negative non weight-observable (contrary to the case of positive weights which is decidable). Beyond active diagnosis of definitive faults, costs and rewards can be introduced in other studies of partially-observed controllable systems: diagnosis when faults are transient, active prediction, opacity preservation/violation, etc.

References

1. N. Bertrand, E. Fabre, S. Haar, S. Haddad, and L. Hélouët. Active diagnosis for probabilistic systems. In *FOSSACS 2014, Grenoble, France*, volume 8412 of *LNCS*, pages 29–42, 2014. doi:10.1007/978-3-642-54830-7_2.
2. C. G. Cassandras and S. Lafortune. *Introduction to Discrete Event Systems - Second Edition*. Springer, 2008.
3. F. Cassez and S. Tripakis. Fault diagnosis with static and dynamic observers. *Fundamenta Informaticae*, 88:497–540, 2008. URL: http://content.iospress.com/articles/fundamenta-informaticae/fi88-4-05.
4. E. Chanthery and Y. Pencolé. Monitoring and active diagnosis for discrete-event systems. In *Proc. SafeProcess'09*, pages 1545–1550, 2009.
5. Stefan Haar, Serge Haddad, Tarek Melliti, and Stefan Schwoon. Optimal constructions for active diagnosis. *J. Comput. Syst. Sci.*, 83(1):101–120, 2017. doi:10.1016/J.JCSS.2016.04.007.
6. M. Sampath, S. Lafortune, and D. Teneketzis. Active diagnosis of discrete-event systems. *IEEE Transactions on Automatic Control*, 43(7):908–929, July 1998. doi:10.1109/9.701089.
7. M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis. Diagnosability of discrete-event systems. *IEEE Trans. Aut. Cont.*, 40(9):1555–1575, 1995. doi:10.1109/9.412626.
8. T-S. Yoo and S. Lafortune. Polynomial-time verification of diagnosability of partially observed discrete-event systems. *IEEE Trans. Automat. Contr.*, 47(9):1491–1495, 2002. doi:10.1109/TAC.2002.802763.
9. Wieslaw Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theor. Comput. Sci.*, 200(1-2):135–183, 1998. doi:10.1016/S0304-3975(98)00009-7.