

Preferential Discrete Model-based Diagnosis for Intermittent and Permanent Faults

Valentin Bouziat¹ and Xavier Pucel¹ and Stéphanie Roussel¹ and Louise Travé-Massuyès²

¹ ONERA / DTIS, Université de Toulouse, F-31055 Toulouse – France

e-mail: firstname.name@onera.fr

²LAAS-CNRS, Université de Toulouse, CNRS, Toulouse, France

e-mail: louise@laas.fr

Abstract

In this paper we consider the diagnosis of intermittent and permanent faults in discrete event systems. We present a logic based modeling approach associated with conditional preferences in order to produce a single diagnosis at each time step. Like all incomplete diagnosis approaches, ours is subject to deadlocks between the system and its diagnoser. In this paper, we address the detection of such deadlocks at design time with the rich semantic model-checker ELECTRUM.

1 Introduction

Using robots in situations where teleoperation is impossible or difficult requires the robots to have some level of decisional autonomy. In particular, an autonomous robot needs to detect and respond appropriately to abnormal losses of performance, due to degradations of the robotic system, or to external disturbances. This requires an elaborate fault management strategy, which can be challenging to design, especially when the mission is complex and the robot is subject to many faults. We are therefore interested in techniques that facilitate the identification of non-nominal operation modes, the detection of faults, and the estimation of the remaining capabilities of the robot.

We propose a modeling formalism constructed so that it meets three requirements. First, it should incrementally calculate a single diagnosis at each time step. This requirement stems from constraints on system performance: one can not afford to compute all possible diagnoses on a large system for a long period of time. Moreover, considering the complete robotic system, the diagnostic module is in our case associated with a deterministic planner which expects a single diagnosis as input. Second, we require our formalism to take into account both permanent and intermittent faults, in particular to distinguish disturbances (noise, false contact, etc.), supposedly intermittent, from degradations (breakdowns, etc.), often permanent. Third, since the diagnosis is used to make autonomous decisions, we want to avoid going back and modifying a diagnosis previously issued. We also forbid from producing a sequence of diagnoses that does not correspond to a possible evolution of the system (for example to remove a permanent fault). Although we do not require strict correctness (we accept delays, slight deviations), this requirement cannot be satisfied if the diagnoser deviates too much from the real system state.

In this paper, we describe how such requirements can create situations where the robotic system produces an observation sequence for which the diagnoser has no explanation. From the point of view of concurrent systems, the physical system sent a message (an observation) that the diagnoser did not expect, thus creating a *deadlock*. In such a situation, the entire decision process supported by the diagnostic engine fails. This is why we are interested in detecting deadlock scenarios at the design time.

The paper is structured as follows. We start with a presentation of the state of the art in Section 2. We then introduce our modeling formalism in Section 3, and describe the deadlock issue in Section 4. We propose a verification method at design time based on the model-checker ELECTRUM [1] in Section 5, and present experimental results on small multi-robot systems in Section 6. Perspectives are discussed in Section 7.

2 Related Work

Our work addresses the diagnosis of discrete event systems. In [2] the authors define a framework for the diagnosis of systems modeled by finite state machines subject to permanent faults. The authors describe the construction of a diagnoser that allows an incremental computation of the diagnosis. However, their diagnoser requires memorizing all the possible state-diagnosis pairs, which severely limits its scalability. Their definition of diagnosability illustrates the importance of validating the performance of a diagnoser at design time.

Our formalism associates propositional logic constraints with a conditional preference theory from [3]. It is inspired by the formalism used in [4] to produce incremental diagnoses. In this related work, the authors just point at the risk of encountering a deadlock. In [5], the same authors propose to detect deadlock scenarios between a system and its diagnoser by an iterative model-checking approach but this approach is difficult to implement and no associated experiments are provided.

The problem of deadlocks between a system and its diagnoser occurs in [6] where the diagnoser goes back in the execution of the diagnoser in order to find a consistent explanation. This solution violates our third requirement explained above. Our goal is to produce reliable diagnoses, or no diagnosis at all.

A classical way to order diagnoses is to prefer the minimal ones, this approach knows several variants compared in [7]. Other approaches like [8] order faults according to some criterion, and deduce an order on diagnoses. In all

these approaches, the diagnosis ordering is unconditional, *i.e.* observations have no effect on the preference order between diagnoses. Our model uses conditional preferences precisely to remove this limitation, and to make it possible to prefer certain diagnoses based on present and past observations.

The diagnosis of intermittent faults is discussed in the literature from different points of view. In [9], the same function is called multiple times, but while faults manifest themselves in an intermittent manner, the underlying diagnosis is constant from one test to another. In [10], a repair event is associated with each fault, and the diagnosis task consists in detecting for each fault which event (fault or repair) happened last. A diagnoser is built to allow incremental evaluation, which involves the previously mentioned scaling problems since all diagnoses are kept in memory.

3 Diagnosis model

Our diagnosis model is a tuple (s_0, Δ, Γ) , where s_0 is the initial system state, Δ is the behavioral model of the system and Γ is the conditional preference model. We assume that the system evolves with discrete events dynamics and that all time steps have the same duration.

3.1 Variables

We use a set of propositional variables P to describe the state of the system. P is partitioned into two subsets O and E , which respectively represent the elements of the system that are observed and those to be estimated. At each discrete time step, given a truth assignment to the variables from O , our goal is to estimate the value for the variables from E .

Notations For a variable set X , we define an *assignment* as a function from X to $\{\top, \perp\}$ that associates to each variable x from X the boolean value *true* (\top) or *false* (\perp). We write x and \bar{x} the assignments to $\{x\}$ such that $x(x) = \top$ and $\bar{x}(x) = \perp$. For a set of variables $X = \{x_1, \dots, x_n\}$, if for $i \in [1, n]$, f_i is an assignment $\{x_i\} \rightarrow \{\top, \perp\}$, then $f_1 f_2 \dots f_n$ denotes the f assignment on X such that $\forall x_i \in X, f(x_i) = f_i(x_i)$. For example, $a \bar{b} \bar{c}$ is the assignment to $\{a, b, c\}$ that assigns true to a and false to b and c .

Definition 1 (Observation). An *observation*, denoted o , is an assignment of the variables of O .

Definition 2 (State). A *state*, denoted s , is an assignment of the variables of P . S denotes the set of states.

3.2 Behavioral Model

The behavioral model $\Delta \subseteq S^2$ is the transition relation of the system. In order to refer to the state of the system at the previous time step, we introduce a bijective function pre on the set P . For a variable p in P , $pre(p)$ represents the value of the variable p at the previous time step. Formally, we define a variable set $P_{pre} = \{pre_p \mid p \in P\}$ such that $\forall p \in P, pre(p) = pre_p$. Δ is represented by a set of propositional logic formulas Δ_p that can relate to both the variables of P and those of P_{pre} .

A pair of states (s_{pre}, s_{now}) belongs to the transition relation Δ when the system can be in the state s_{pre} at time $t - 1$ and in the state s_{now} at time t . To formally link Δ to Δ_p , for any pair of states (s_{pre}, s_{now}) , we define the assignment $\sigma_{s_{pre}, s_{now}}$ on variables from $P \cup P_{pre}$ such that $\forall p \in P, \sigma_{s_{pre}, s_{now}}(p) = s_{now}(p)$ and $\sigma_{s_{pre}, s_{now}}(pre(p)) =$




Figure 1: A simple system composed of a lamp and a switch

$s_{pre}(p)$. We consider that a pair of states belongs to the transition relation Δ if and only if the corresponding assignment satisfies the formulas of Δ_p .

Definition 3 (Transition). A pair of states $(s_{pre}, s_{now}) \in S^2$ is a *transition*, denoted $(s_{pre}, s_{now}) \in \Delta$, if and only if $\sigma_{s_{pre}, s_{now}} \models \Delta_p$.

In the remaining of this paper, we do not distinguish Δ and Δ_p .

In order to reason about the possible evolutions of the system state, and the associated observation sequences, we define consistent state sequences and consistent observation sequences as follows.

Definition 4 (Consistent state sequence). A *state sequence* $(s_0, s_1, \dots, s_n) \in S^n$ is *consistent* if and only if $\forall i \in [1, n], (s_{i-1}, s_i) \in \Delta$.

Definition 5 (Consistent observation sequence). An *observation sequence* (o_0, o_1, \dots, o_n) is *consistent* if and only if there exist a consistent state sequence (s_0, s_1, \dots, s_n) such that $\forall i \in [0, n], \forall o \in O, s_i(o) = o_i(o)$.

In the following, when there is no ambiguity, state and observation sequences are assumed to be consistent.

Given a previous state and an observation of the system, we define the set of candidates for diagnosis as the set of states that are compatible with the previous state, the observation and the behavioral model Δ .

Definition 6 (Diagnosis candidates). The set $S_\Delta(s_{pre}, o)$ of *diagnosis candidates* for a previous state s_{pre} and an observation o is defined by:

$$S_\Delta(s_{pre}, o) = \{s_{now} \in S \mid (s_{pre}, s_{now}) \in \Delta \text{ and } \forall o \in O, s_{now}(o) = o(o)\}$$

Example 1. The system illustrated in Figure 1 is composed of a lamp controlled by a switch. This system may be subject to a permanent fault and an intermittent fault, our goal is to estimate their presence or absence. The set O is composed of $light$ which is true when the light is on, false otherwise, and $switch$ which is true when the switch is closed (current flows), false otherwise. The set E contains two variables representing the two faults that may occur: f_{wire} represents an intermittent loose contact in the wire and f_{light} a permanent lamp failure.

The two rules of Δ represent the following behavior: the lamp glows when the switch is closed and when there is no fault on the wire nor in the lamp (δ_1). If the bulb of the lamp was broken at the previous state, then it is at the present time (δ_2), *i.e.* the fault f_{light} is permanent. In the initial state $s_0 = light \ switch \ f_{light} \ f_{wire}$, the lamp glows and there is no fault.

$$\Delta = \left\{ \begin{array}{l} \text{light} \leftrightarrow \text{switch} \wedge \neg \text{f}_{\text{light}} \wedge \neg \text{f}_{\text{wire}} \\ \text{pre_f}_{\text{light}} \rightarrow \text{f}_{\text{light}} \end{array} \begin{array}{l} (\delta_1) \\ (\delta_2) \end{array} \right\}$$

From the behavioral model Δ , we can calculate diagnosis candidates. For instance, at time step 1, s_0 is the previous system state, let us consider the observation $o_1 = \overline{\text{light switch}}$. The candidate states set is $\mathcal{S}_\Delta(s_0, o_1) = \{s_{\text{now}}, s'_{\text{now}}, s''_{\text{now}}\}$ with:

$$\begin{aligned} s_{\text{now}} &= \overline{\text{light switch}} \overline{\text{f}_{\text{light}}} \overline{\text{f}_{\text{wire}}} \\ s'_{\text{now}} &= \overline{\text{light switch}} \text{f}_{\text{light}} \overline{\text{f}_{\text{wire}}} \\ s''_{\text{now}} &= \overline{\text{light switch}} \text{f}_{\text{light}} \text{f}_{\text{wire}} \end{aligned}$$

In the state s_{now} , the intermittent fault is present alone; in s'_{now} , the permanent fault is present alone; both faults are present in s''_{now} .

3.3 Single diagnosis choice

For a given previous state and observation, in order to produce a single diagnosis, we have to choose a single state from all the candidates of $\mathcal{S}_\Delta(s_{\text{pre}}, o)$. This choice is dictated by the conditional preference model Γ that consists in an ordered set of *conditional preferences*.

A conditional preference relates to a variable to be estimated and indicates under which condition we prefer the diagnoses that assign true or false to this variable to the other diagnoses. Our modeling formalism makes it easy for a preference condition to refer to past variables, which would not be possible with the more usual “next” operator [11]. A preference is a form of “soft constraint”, it is only applied when there exists diagnoses with both values for the variable. A formal definition follows.

Definition 7 (Conditional preference). A conditional preference γ on a variable e of \mathbf{E} , is denoted $\text{cond} : e \prec \bar{e}$. The preference’s condition cond is a propositional formula on $\mathbf{P} \cup \mathbf{P}_{\text{pre}}$. The variable e is called the preference’s target.

Note that $\text{cond} : \bar{e} \prec e$ is equivalent to $\neg \text{cond} : e \prec \bar{e}$. For a variable e from \mathbf{E} , the conditional preference $\text{cond} : e \prec \bar{e}$ expresses a preference for states in which e is true to those where e is false if and only if cond is satisfied.

Formally, a preference $\gamma = \text{cond} : e \prec \bar{e}$ defines a partial order \prec_γ an equivalence relation \approx_γ between pairs of transitions as follows. For all triples $s_{\text{pre}}, s, s' \in \mathcal{S}^3$, γ strictly prefers the transition (s_{pre}, s) to transition (s_{pre}, s') (denoted $(s_{\text{pre}}, s) \prec_\gamma (s_{\text{pre}}, s')$) if and only if $\sigma_{s_{\text{pre}}, s} \models \text{cond} \leftrightarrow e$ and $\sigma_{s_{\text{pre}}, s'} \not\models \text{cond} \leftrightarrow e$. Transitions (s_{pre}, s) and (s_{pre}, s') are equivalent to γ (denoted $(s_{\text{pre}}, s) \approx_\gamma (s_{\text{pre}}, s')$) if and only if $(\sigma_{s_{\text{pre}}, s} \models \text{cond} \leftrightarrow e) \Leftrightarrow (\sigma_{s_{\text{pre}}, s'} \models \text{cond} \leftrightarrow e)$.

Example 2. Let us consider the preference $\gamma = \neg \text{light} : \text{f}_{\text{light}} \prec \overline{\text{f}_{\text{light}}}$. This preference indicates that if both f_{light} and $\overline{\text{f}_{\text{light}}}$ are part of some diagnosis candidate, then f_{light} is preferred if and only if $\neg \text{light}$ holds. Formally, we prefer diagnoses that satisfy $\neg \text{light} \leftrightarrow \text{f}_{\text{light}}$. For the states $s_{\text{now}}, s'_{\text{now}}$ and s''_{now} in Example 1, as $\neg \text{light}$ holds, we prefer the states in which f_{light} holds, i.e. the states s'_{now} and s''_{now} . Formally, $s'_{\text{now}} \prec_\gamma s_{\text{now}}$, $s''_{\text{now}} \prec_\gamma s_{\text{now}}$ and $s''_{\text{now}} \approx_\gamma s'_{\text{now}}$.

We assume that in Γ , all estimated variables are the target of a conditional preference. While not always necessary, it

ensures the diagnoser is deterministic. This raises the question of the order in which preferences are applied, that we address now.

Definition 8 (Conditional preference model). A conditional preference model Γ is a sequence of conditional preferences $(\gamma_1, \gamma_2, \dots, \gamma_n)$ with $\gamma_i = \text{cond}_i : e_i \prec \bar{e}_i$ for $i \in [1, n]$ such that each variable e from \mathbf{E} is the target of exactly one preference.

We require Γ to be an acyclic preference model, which guarantees its consistency (see [12]). For example, the following two preferences form a cycle: $a : b \prec \bar{b}$ and $b : a \prec \bar{a}$. In the first preference the value of a depends on that of b , and in the second preference the exact opposite happens. While the literature contains work on cyclic preference networks [12], in this paper we assume that Γ is acyclic. This means that the condition of a preference γ_i cannot use variables that are the target of the following preferences in the sequence. Formally, $\forall i \in [1, n]$, the scope of the condition cond_i is a subset of $\mathbf{P}_{\text{pre}} \cup \mathbf{O} \cup \{e_j \mid 1 \leq j < i\}$.

From a preference model Γ , it is possible to define a partial order \prec_Γ between pairs of states as follows. Intuitively, as in a lexicographic order, we consider preferences γ_i in their index order in Γ and we apply at each index the order relation \prec_{γ_i} defined above. This order is partial because it compares only pairs of transitions $(s_{\text{pre}}, s_{\text{now}})$ and $(s_{\text{pre}}', s_{\text{now}}')$ that have the same previous state (i.e. $s_{\text{pre}} = s_{\text{pre}}'$) and whose successor states produce the same observation (i.e. $\forall o \in \mathbf{O} s_{\text{now}}(o) = s_{\text{now}}'(o)$).

Formally, $\forall s_{\text{pre}}, s, s' \in \mathcal{S}^3$, (s_{pre}, s) is strictly preferred to (s_{pre}, s') by Γ (denoted $(s_{\text{pre}}, s) \prec_\Gamma (s_{\text{pre}}, s')$) if and only if there exists $i \in [1, n]$ such that for all $j < i$, $(s_{\text{pre}}, s) \approx_{\gamma_j} (s_{\text{pre}}, s')$ and $(s_{\text{pre}}, s) \prec_{\gamma_i} (s_{\text{pre}}, s')$.

Although the \prec_Γ order is partial on \mathcal{S}^2 , it is complete on any set of diagnosis candidates. We use this order relation to define the preferred diagnosis at each time step.

Proposition 1. Let s_{pre} be a state, o an observation, Δ a behavioral model and Γ a preference model. If s and s' are two candidate states for s_{pre} and o ($s, s' \in \mathcal{S}_\Delta(s_{\text{pre}}, o)^2$) such that $s \neq s'$ then we have either $(s_{\text{pre}}, s) \prec_\Gamma (s_{\text{pre}}, s')$ or $(s_{\text{pre}}, s') \prec_\Gamma (s_{\text{pre}}, s)$.

Proof By Definition 6, s and s' belong to $\mathcal{S}_\Delta(s_{\text{pre}}, o)$ implies that $\forall o \in \mathbf{O}, s(o) = s'(o)$. Therefore, $s \neq s'$ means that there is a variable $e \in \mathbf{E}$ such that $s(e) \neq s'(e)$. So, there exists a preference $\gamma \in \Gamma$ such that $s \approx_\gamma s'$ does not hold. Let i be the index of the first preference of the sequence in this case. Formally, γ_i is such that $\forall j < i, (s_{\text{pre}}, s) \approx_{\gamma_j} (s_{\text{pre}}, s')$, $(s_{\text{pre}}, s) \approx_{\gamma_i} (s_{\text{pre}}, s')$ does not hold. This means that $(s_{\text{pre}}, s) \prec_{\gamma_i} (s_{\text{pre}}, s')$ or $(s_{\text{pre}}, s') \prec_{\gamma_i} (s_{\text{pre}}, s)$. From the order \prec_Γ , we then have $(s_{\text{pre}}, s) \prec_\Gamma (s_{\text{pre}}, s')$ or $(s_{\text{pre}}, s') \prec_\Gamma (s_{\text{pre}}, s)$. \square

3.4 Estimation process

At each step, given the set of candidate states $\mathcal{S}_\Delta(s_{\text{pre}}, o)$ we select the state preferred by the conditional preference model Γ .

Definition 9 (Estimated state). The estimated state $\hat{s} = \text{estim}(s_{\text{pre}}, o)$ for a previous state s_{pre} and an observation o is the element of $\mathcal{S}_\Delta(s_{\text{pre}}, o)$ preferred by \prec_Γ . Formally, $\hat{s} = \text{estim}(s_{\text{pre}}, o)$ if and only if:

1. $\hat{s} \in \mathcal{S}_\Delta(s_{\text{pre}}, o)$, and

2. $\forall s_{now} \in \mathcal{S}_\Delta(s_{pre}, o)$ such that $s_{now} \neq \hat{s}$, we have $(s_{pre}, \hat{s}) \prec_\Gamma (s_{pre}, s_{now})$.

From the initial state s_0 , we now define the sequence of estimated states that is produced by the diagnoser for a given observation sequence.

Definition 10 (Estimated state sequence). *A state sequence $(s_0, \hat{s}_1, \hat{s}_2, \dots, \hat{s}_k)$ is the estimated state sequence for the observation sequence $(o_0, o_1, o_2, \dots, o_k)$ if and only if $\forall i \in [1, k]$, $\hat{s}_i = estim(\hat{s}_{i-1}, o_i)$.*

Example 3. *Let us associate the lamp model Δ from Example 1 to the following preferences:*

$$\Gamma = \left(\begin{array}{l} \text{pre_f}_{\text{wire}} : \overline{f_{\text{wire}}} \prec \overline{f_{\text{wire}}} \quad (\gamma_1) \\ \perp : f_{\text{light}} \prec \overline{f_{\text{light}}} \quad (\gamma_2) \end{array} \right)$$

The first preference (γ_1) declares that we prefer the value for f_{wire} that was estimated at the previous time step. This mechanism makes it possible to bring some stability to the diagnosis since f_{wire} is an intermittent fault: if Δ allows both diagnoses for f_{wire} , we prefer to maintain the diagnosis that was chosen at the previous time step.

The preference (γ_2) indicates that we always prefer to assume that f_{light} is absent.

In Example 1, with $s_0 = \overline{\text{light switch}} \overline{f_{\text{light}}} \overline{f_{\text{wire}}}$ as the previous state and $o_1 = \overline{\text{light switch}}$ as the observation, the set of diagnosis candidates $\mathcal{S}_\Delta(s_{pre}, o)$ contains the three states $s_{now} = \overline{\text{light switch}} \overline{f_{\text{light}}} \overline{f_{\text{wire}}}$, $s_{now}' = \overline{\text{light switch}} \overline{f_{\text{light}}} f_{\text{wire}}$ and $s_{now}'' = \overline{\text{light switch}} f_{\text{light}} \overline{f_{\text{wire}}}$.

By applying (γ_1) first, as $\text{pre_f}_{\text{wire}}$ holds in the three states, we prefer the states in which $\overline{f_{\text{wire}}}$. In our case, only s_{now}' satisfies this criterion. As there is only one diagnosis left, we do not need to apply preference (γ_2) . The preferred single state is $\hat{s}_1 = estim(s_0, o_1) = s_{now}' = \overline{\text{light switch}} \overline{f_{\text{light}}} \overline{f_{\text{wire}}}$.

4 Deadlock

A deadlock situation occurs when the diagnoser has previously estimated a state \hat{s} different from the current system state s , and receives an observation o in contradiction with \hat{s} and Δ . In such a situation the candidates set for \hat{s} and o is empty and the diagnoser is unable to return a diagnosis.

Definition 11 (Deadlock). *An estimation model (s_0, Δ, Γ) is in a deadlock situation for an observation sequence $(o_0, o_1, o_2, \dots, o_k)$ with $k > 1$ ¹ if and only if:*

1. there exists an estimated state sequence for (o_0, \dots, o_{k-1}) , and
2. there exist no estimated state sequence for (o_0, \dots, o_k)

Example 4. *Let (o_0, o_1, o_2, o_3) be the observation sequence produced by the system and along with the associated state sequence (s_0, s_1, s_2, s_3) described in Figure 2.*

When observation o_1 is received, Γ selects the preferred state $\hat{s}_1 = \overline{\text{light switch}} \overline{f_{\text{light}}} \overline{f_{\text{wire}}}$. Then for observation o_2 , we are in the situation described in Example 3 and the preferred state is $\hat{s}_2 = \overline{\text{light switch}} \overline{f_{\text{light}}} \overline{f_{\text{wire}}}$. Observation o_3 is in contradiction with the previously estimated state. In fact the estimator has estimated f_{light} in the previous state, meaning that variable $\text{pre_f}_{\text{light}}$ is true. Rules of Δ are not consistent with such a configuration: (δ_1) requires f_{light} to

¹There can be no deadlock at \hat{s}_1 because we assume that the diagnoser is correctly initialized at s_0 .

be false because the lamp glows while (δ_2) requires f_{light} to be true since the associated fault is permanent.

Therefore, there is an estimated state sequence for (o_0, o_1, o_2) : this is the sequence $(s_0, \hat{s}_1, \hat{s}_2)$. However, since there is no estimated state sequence for (o_0, o_1, o_2, o_3) , the pair system-diagnoser is in a deadlock situation for this observation sequence.

In the previous example, the deadlock situation could be eliminated by changing the order or conditions of preferences. However, as soon as one is interested in more complex systems, it becomes difficult to anticipate the deadlock situations and even more to solve them. In this paper, we focus on detecting these deadlocks and leave their resolution for future work.

5 Deadlock Checking

In this section, we show how to use a model-checker to identify deadlock scenarios at the design phase of the diagnostic model.

5.1 Deadlock checker

The verification method we propose is inspired from the Twin-Plant [13] technique, which makes it possible to verify the diagnosability of a system by constructing the synchronous product of two finite state machines.

Our method is similar since it involves constructing a deadlock verifier by synchronizing two state machines. The first state machine represents the system and is constrained only by the transition relation Δ . We use it to generate consistent observation sequences. The second state machine is the diagnoser built from the whole estimation model. It is also constrained by Δ , but it deterministically selects the next state by applying Γ preferences. The verifier is the product of these two state machines synchronized on observable variables at each time step. Deadlock checking then consists in checking whether there exists an observation sequence that leads the verifier to a state in which the system has a successor state, but the diagnoser has none.

In order to distinguish the states of the two state machines presented above, *i.e.* the state of the system and that of the diagnoser, we introduce two variable sets P^{sys} and P^{est} that are direct copies of P . We also use a est_sat variable that indicates whether the diagnoser has a successor state.

Definition 12 (Verifier variables). *The set of the verifier variables is defined by $\text{P}^{\text{verif}} = \text{P}^{\text{sys}} \cup \text{P}^{\text{est}} \cup \{\text{est_sat}\}$, where :*

- $\text{P}^{\text{sys}} = \{\text{p_sys} \mid \text{p} \in \text{P}\}$ is the set of variables describing the system state ;
- $\text{P}^{\text{est}} = \{\text{p_est} \mid \text{p} \in \text{P}\}$ is the set of variables describing the diagnoser state;

Step i	o_i	s_i	\hat{s}_i
0	$\overline{\text{light switch}}$	$\overline{f_{\text{light}}} \overline{f_{\text{wire}}}$	$\overline{f_{\text{light}}} \overline{f_{\text{wire}}}$
1	$\overline{\text{light switch}}$	$\overline{f_{\text{light}}} \overline{f_{\text{wire}}}$	$\overline{f_{\text{light}}} \overline{f_{\text{wire}}}$
2	$\overline{\text{light switch}}$	$\overline{f_{\text{light}}} \overline{f_{\text{wire}}}$	$\overline{f_{\text{light}}} \overline{f_{\text{wire}}}$
3	$\overline{\text{light switch}}$	$\overline{f_{\text{light}}} \overline{f_{\text{wire}}}$	/

Figure 2: Deadlock scenario for Example 4. In columns s_i and \hat{s}_i , we omit variables from 0 whose values are identical to the column o_i , and we only represent variables from E).

- `est_sat` indicates whether there is an estimated state for the diagnoser.

A state of the verifier is an assignment on variables of $\mathcal{P}^{\text{verif}}$.

In order to define the state machine resulting from the synchronous product, we successively define the initial state of the verifier (Definition 13) and its transition relation (Definition 14).

Definition 13 (Verifier initial state). *The initial state s_0^{verif} of the verifier is such that:*

- $\forall p \in \mathcal{P}, s_0^{\text{verif}}(\text{p_sys}) = s_0(p)$,
- $\forall p \in \mathcal{P}, s_0^{\text{verif}}(\text{p_est}) = s_0(p)$,
- $s_0^{\text{verif}}(\text{est_sat}) = \top$.

Definition 14 (Verifier transition relation). *Let $s_{\text{pre}}^{\text{verif}}$ and $s_{\text{now}}^{\text{verif}}$ be two verifier states. The pair $(s_{\text{pre}}^{\text{verif}}, s_{\text{now}}^{\text{verif}})$ is a verifier transition if and only if:*

- (1) variables associated with observations take the same value on system and diagnoser sides,
- (2) the system transition described by variables of \mathcal{P}^{sys} satisfies Δ ,
- (3) the variable `est_sat` indicates whether there exists a possible estimated state (i.e. if the candidates set is not empty), and
- (4) if `est_sat` is true, then the diagnoser transition described by variables of \mathcal{P}^{est} satisfies Δ and Γ .

Formally :

$$\forall o \in \mathcal{O}, s_{\text{pre}}^{\text{verif}}(\text{sys_o}) = s_{\text{pre}}^{\text{verif}}(\text{est_o}) \text{ and} \quad (1)$$

$$\forall o \in \mathcal{O}, s_{\text{now}}^{\text{verif}}(\text{sys_o}) = s_{\text{now}}^{\text{verif}}(\text{est_o})$$

$$\exists (s_{\text{pre}}, s_{\text{now}}) \in \Delta, \forall p \in \mathcal{P},$$

$$s_{\text{pre}}(p) = s_{\text{pre}}^{\text{verif}}(\text{sys_p}) \text{ and} \quad (2)$$

$$s_{\text{now}}(p) = s_{\text{now}}^{\text{verif}}(\text{sys_p})$$

$$\text{est_sat} \leftrightarrow$$

$$\left(\begin{array}{l} \exists (s_{\text{pre}}, s_{\text{now}}) \in \Delta, \\ \forall p \in \mathcal{P}, s_{\text{pre}}(p) = s_{\text{pre}}^{\text{verif}}(\text{est_p}) \text{ and} \\ \forall p \in \mathcal{O}, s_{\text{now}}(p) = s_{\text{now}}^{\text{verif}}(\text{est_p}) \end{array} \right) \quad (3)$$

$$\text{est_sat} \rightarrow$$

$$\left(\begin{array}{l} \exists (s_{\text{pre}}, s_{\text{now}}) \in \Delta, \forall p \in \mathcal{P}, \\ s_{\text{pre}}(p) = s_{\text{pre}}^{\text{verif}}(\text{est_p}) \text{ and} \\ s_{\text{now}}(p) = s_{\text{now}}^{\text{verif}}(\text{est_p}) \text{ and} \\ \left(\begin{array}{l} \exists s_{\text{best}} \in \mathcal{S}, \\ \forall o \in \mathcal{O}, s_{\text{best}}(o) = s_{\text{now}}(o), \text{ and} \\ (s_{\text{pre}}, s_{\text{best}}) \in \Delta, \text{ and} \\ (s_{\text{pre}}, s_{\text{best}}) \prec_{\Gamma} (s_{\text{pre}}, s_{\text{now}}) \end{array} \right) \end{array} \right) \quad (4)$$

Proposition 2. *An estimation model (s_0, Δ, Γ) is subject to a deadlock if and only if the associated verifier contains a path in which `est_sat` is false. The deadlock scenario is the observation sequence corresponding to the states of the verifier.*

Proof Suppose that the estimation model is subject to a deadlock. According to Definition 11, there exists an observation sequence (o_0, o_1, \dots, o_n) generated by a consistent state sequence (s_0, s_1, \dots, s_n) , such that there is an

estimated state sequence $(\hat{s}_0, \hat{s}_1, \dots, \hat{s}_{n-1})$ for the partial sequence $(o_0, o_1, \dots, o_{n-1})$ and that there does not exist an estimated state sequence $(\hat{s}_0, \hat{s}_1, \dots, \hat{s}_n)$ for the complete sequence. We then build the verifier's state sequence $(s_0^{\text{verif}}, s_1^{\text{verif}}, \dots, s_n^{\text{verif}})$ as follows:

- $\forall i \in [0, n], \forall p \in \mathcal{P}, s_i^{\text{verif}}(\text{p_sys}) = s_i(p)$,
- $\forall i \in [0, n-1], \forall p \in \mathcal{P}, s_i^{\text{verif}}(\text{p_est}) = \hat{s}_i(p)$,
- $\forall i \in [0, n-1], s_i^{\text{verif}}(\text{est_sat}) = \top$,
- $\forall p \in \mathcal{O}, s_n^{\text{verif}}(\text{p_est}) = o_n(p)$,
- $\forall p \in \mathcal{E}, s_n^{\text{verif}}(\text{p_est}) = \top$,
- $s_n^{\text{verif}}(\text{est_sat}) = \perp$.

We show that for all $i \in [0, n-1]$, $(s_i^{\text{verif}}, s_{i+1}^{\text{verif}})$ is a transition of the checker: let i be an integer in $[0, n-1]$, we show that the four parts of Definition 14 are satisfied:

- $\forall o \in \mathcal{O}, \forall p \in \mathcal{P}, s_i(p) = \hat{s}_i(p)$ and $s_{i+1}(p) = \hat{s}_{i+1}(p)$. Equation (1) holds;
- equation (2) holds by construction of the verifier;
- for $i < n-1$, $s_i^{\text{verif}}(\text{est_sat}) = \top$ and $(\hat{s}_i, \hat{s}_{i+1})$ is the pair of states in Δ satisfying equation (3);
- for $i = n-1$, the definition of a deadlock implies that there does not exist a pair of states satisfying Δ : equation (3) is satisfied;
- for $i < n-1$, $s_i^{\text{verif}}(\text{est_sat}) = \top$ and $(\hat{s}_i, \hat{s}_{i+1})$ is the pair of states in Δ satisfying equation (4);

For the reciprocal, let us assume that the checker has a path in which `est_sat` is false. We consider such a path $(s_0^{\text{verif}}, s_1^{\text{verif}}, \dots, s_n^{\text{verif}})$ in which s_n^{verif} is the only state in which false is assigned to `est_sat`. We then build two states sequences (s_0, s_1, \dots, s_n) and $(s_0, \hat{s}_1, \dots, \hat{s}_{n-1})$ as follows:

- $\forall i \in [0, n], \forall p \in \mathcal{P}, s_i(p) = s_i^{\text{verif}}(\text{p_sys})$,
- $\forall i \in [0, n-1], \forall p \in \mathcal{P}, \hat{s}_i(p) = s_i^{\text{verif}}(\text{p_est})$

We prove that these state sequences are consistent with Δ . From Definition 14, `est_sat` is false (last time step) implies that no state s_{now} allows to satisfy the right part of the condition (equation 3) for $s_{\text{pre}} = \hat{s}_{n-1}$. This means that no state meets the Definition 6. There is a deadlock for the observation sequence generated by (s_0, s_1, \dots, s_n) . \square

5.2 Choice of the model-checker

Model-checking is a set of techniques and computer tools for checking properties on systems. Among the many model-checkers in the literature, NuSMV [14] and NuXMV are known for their effectiveness in checking temporal properties on dynamic systems. They are based on temporal logic like LTL or CTL. Another family of model checkers is specialized in checking properties on structurally complex models, but without temporal dynamics. This is the case of the Alloy Analyzer model-checker [15] which supports a fragment of first order logic.

In the case of deadlock verification, we need to express the temporal dynamics of the system but also the preferences of the diagnoser, which are not easily expressed in propositional logic. We therefore chose to use the model-checker ELECTRUM [1] which is based on the Alloy language and integrates temporal dynamics as in NuSMV.

5.3 Quantified form of preference

The concept of preference is part of the verifier definition. This means that the model-checker language must allow to express them. To do so, we define a *quantified form* for preferences.

Definition 15 (Preference quantified form). *Let be $\Gamma = (\gamma_1, \dots, \gamma_n)$ a preference model in which each preference has form $\gamma_i = \text{cond}_i : e_i \prec \bar{e}_i$. The quantified form of γ_i is the formula ψ_i defined by:*

$$\psi_i = (\forall e_i, \exists e_{i+1}, \dots, e_n, \Delta) \rightarrow (e_i \leftrightarrow \text{cond}_i)$$

The quantified form ψ_i is composed of two parts. The left part is only satisfied when for both truth values variable e_i , there is a way to assign the target variables of the following preferences and still satisfy Δ . It represents the cases when the preference is actually applied. The right part expresses the effect of the preference, *i.e.* that e_i is assigned to true if and only if the condition cond_i is satisfied. Through the following proposition, we show that Γ preferences and their quantified form are equivalent from the point of view of the estimated state for a previous state and an observation.

Proposition 3. *For a previous estimated state \hat{s}_{pre} , an observation o , a candidate $\hat{s} \in \mathcal{S}_\Delta(\hat{s}_{pre}, o)$, $\hat{s} = \text{estim}(\hat{s}_{pre}, o)$ if and only if the assignment $\sigma_{\hat{s}_{pre}, \hat{s}}$ satisfies conjunction $\bigwedge_{i \in [1, n]} \psi_i$.*

Proof. Let \hat{s}_{pre} be a previous estimated state, o an observation and \hat{s} a state in $\mathcal{S}_\Delta(\hat{s}_{pre}, o)$. We inductively show that \hat{s} is preferred for the preference sequence $(\gamma_1, \dots, \gamma_k)$ if and only if the assignment $\sigma_{\hat{s}_{pre}, \hat{s}}$ satisfies $\bigwedge_{i \in [1, k]} \psi_i$.

For $k = 1$, the set of free variables in formula $lhs_1 = \forall e_1, \exists e_2, \dots, e_n, \Delta$ is the set \emptyset , since all the variables in E are quantified. The formula lhs_1 is satisfied by an observation o if and only if both oe_1 and $o\bar{e}_1$ have an extension in $P \cup P_{pre}$ that satisfies Δ . This means $\mathcal{S}_\Delta(\hat{s}_{pre}, o)$ contains at least two states s and s' with $s(e_1) \neq s'(e_1)$. Among these two states, the one that satisfies $\text{cond}_1 \leftrightarrow e_1$ is the preferred state by preference γ_1 , and also satisfies ψ_1 . As the states preferred by γ_1 are exactly the ones for which the assignment $\sigma_{\hat{s}_{pre}, \hat{s}}$ satisfies ψ_1 , then the proposition holds for $k = 1$.

Let $k > 1$ and let us assume that the induction is valid at index $k - 1$, *i.e.* that \hat{s} is preferred for the preference sequence $(\gamma_1, \dots, \gamma_{k-1})$ if and only if the assignment $\sigma_{\hat{s}_{pre}, \hat{s}}$ satisfies $\bigwedge_{i \in [1, k-1]} \psi_i$. Since s and s' belong to $\mathcal{S}_\Delta(\hat{s}_{pre}, o)$, since s is a state preferred for $(\gamma_1, \dots, \gamma_{k-1})$ and since $\sigma_{\hat{s}_{pre}, s'}$ satisfies $\bigwedge_{i \in [1, k-1]} \psi_i$, then for all $i < k$, $s(e_i) = s'(e_i)$.

The set of free variables in formula $lhs_k = \forall e_k, \exists e_{k+1}, \dots, e_n, \Delta$ is the set $\emptyset \cup \{e_i \mid i < k\}$. An assignment op to these variables satisfies lhs_k if and only if $op e_k$ and $op \bar{e}_k$ both have an extension to $P \cup P_{pre}$ that satisfies Δ .

If for a given assignment op to $\emptyset \cup \{e_i \mid i < k\}$ the formula lhs_k does not hold, then for all s and s' in $\mathcal{S}_\Delta(\hat{s}_{pre}, o)$, $s(e_k) = s'(e_k)$. Thus, the preference γ_k is not applied and the preferred states for $(\gamma_1, \dots, \gamma_{k-1})$ are the also preferred for $(\gamma_1, \dots, \gamma_k)$. In this case, $\psi_k = \top$ and the states s' such that $\sigma_{\hat{s}_{pre}, s'}$ satisfy $\bigwedge_{i \in [1, k-1]} \psi_i$ are the same ones for which $\sigma_{\hat{s}_{pre}, s'}$ satisfy $\bigwedge_{i \in [1, k]} \psi_i$ and the induction is valid.

If for a given assignment op to $\emptyset \cup \{e_i \mid i < k\}$ the formula lhs_k holds, then $\mathcal{S}_\Delta(\hat{s}_{pre}, o)$ contains at least two states s

```

pred delta[fwire, light, flight,
           switch, preF_flight : Bool] {
  (isTrue[preF_flight] => isTrue[flight])
  and
  (isTrue[light] <=> ( isTrue[switch] and
                       !isTrue[flight] and
                       !isTrue[fwire]))
}
var one sig RealSystem {
  var fwire : Bool,
  var preF_fwire : Bool,
  var light : Bool,
  var flight : Bool,
  var switch : Bool,
  var preF_flight : Bool,
}
fact always_delta_RealSystem {
  always {
    delta[RealSystem.fwire,
          RealSystem.preF_fwire,
          RealSystem.light,
          RealSystem.flight,
          RealSystem.switch,
          RealSystem.preF_flight]}
}
fact synch_obs {
  always {
    Estimator.light = RealSystem.light
    Estimator.switch = RealSystem.switch}
}

```

Figure 3: The observed system and observation synchronization in ELECTRUM.

and s' such that for all $i < k$, $s(e_i) = s'(e_i)$ and $s(e_k) \neq s'(e_k)$. Between these two states, the one in which $\text{cond}_k \leftrightarrow e_k$ is true is the preferred state by preference γ_k . As the states preferred by γ_k are exactly the ones for which the assignment $\sigma_{\hat{s}_{pre}, \hat{s}}$ satisfies ψ_k , thus the induction is valid at index k . \square

5.4 Model Encoding with ELECTRUM

In this section, we illustrate how the verifier defined in Definition 13 and 14 is encoded in ELECTRUM. We illustrate the encoding of the system described in Examples 1 and 3.

To model the system's state machine, we declare a predicate corresponding exactly to the model Δ then we declare a constraint (a `fact` in ELECTRUM) that specifies that the system respects the predicate at all time steps (`always` in ELECTRUM), as illustrated in Figure 3. A similar constraint is declared for the diagnoser's state machine. Finally, we synchronize the variables from \emptyset of the observed system with those of the second state machine representing our estimator, as expressed in condition (1) of Definition 14.

To represent the temporal aspect of our formalism, we use the operator `'` from ELECTRUM which describes a variable at the next time step. Figure 4 reproduces the intention of the bijective function `pre` (`p_pre` value at next state is equal to `p` value at current state).

ELECTRUM supports the use of relational algebra operators. It is possible for each preference to write a predicate telling us if a preference γ_i is applicable, that is, if the two valuations for variable e_i (operator `all`), and at least one possible valuation for variables $e_j, i < j \leq n$ (operator `some`) are consistent with Δ an the observation. This corresponds to the left part of the quantified form of a preference. Figure 5 illustrates the definition of predicates that implement these conditions with ELECTRUM operators. The rest

```

fact next_Estimator {
  always {
    Estimator.preF_flight' =
      Estimator.flight
    Estimator.preF_fwire' =
      Estimator.fwire }
}

```

Figure 4: Temporal dynamics in ELECTRUM.

```

pred pref_fwire_possible{
  all fwire : Bool | some flight : Bool |
  delta[fwire,
  Estimator.preF_fwire,
  Estimator.light,
  flight,
  Estimator.switch,
  Estimator.preF_flight]
}
pred pref_fwire_applied{
  isTrue[Estimator.fwire] <=>
  isTrue[Estimator.preF_fwire]
}
pred pref_flight_possible{
  all flight : Bool |
  delta[Estimator.fwire,
  Estimator.preF_fwire,
  Estimator.light,
  flight,
  Estimator.switch,
  Estimator.preF_flight]
}
pred pref_flight_applied{
  isTrue[Estimator.flight] <=>
  isTrue[False]
}
fact prefs {
  always {
    pref_flight_possible implies
      pref_flight_applied
    pref_fwire_possible implies
      pref_fwire_applied
  }
}

```

Figure 5: Implementing preferences in ELECTRUM

of the preferences are described by a fact.

The order in which preferences are applied in Γ , is completely induced by the quantification overlays in the quantified forms of preferences (see Definitions 8 and 15).

6 Experimental results

We experimented our approach on the model of a multi-robot mission of customizable size and complexity.

In this mission, one or more robots move on a rectangular grid of variable size, each robot moves towards its destination. The destination can change during the mission according to unspecified dynamics.

Behavioral constraints are declared in Δ for each robot. robots should move unless they are at their destination. Moreover, robots may be subject to permanent and / or intermittent faults. An intermittent fault slows down or immobilizes the robot for a few moments which is modelled by the fact that a robot takes several time steps to cross a cell. A permanent fault immobilizes the robot permanently. We also estimate the state of the terrain. A robot crosses a *normal* cell in one time step, but we introduce *difficult* cells

(the robot takes several time steps to go through the cell) and *dangerous* cells (robot stays forever on the cell).

We observe at every moment the position of robots on the grid and their destination. We estimate the presence of intermittent and permanent faults on each robot, a variable indicating if it can move, as well as the dangerousness of each grid cell. We implement two estimation strategies, one subject to deadlock, the other not (for example always preferring the absence of a permanent fault).

Models are written in Scala language and are automatically translated into ELECTRUM. ELECTRUM proceeds to the verification through different solvers. For our experiments, we set ELECTRUM to use the Sat4j solver [16]. Verifications were performed on a 3.5GHz Intel Core i5-7600 quad-core processor. Figure 6 shows the results of our experiments for different mission parameters.

The first conclusion is that verification is faster for models that are subject to deadlock scenarios than for those that do not encounter such scenarios. This difference is directly related to the fact that when the model-checker finds a deadlock scenario in the model, the search ends immediately and it returns a counterexample corresponding to the sequence of observations. On the contrary, to prove that a model contains no deadlock scenarios, ELECTRUM explores many more possible executions paths for the verifier.

We can then see that the verification is very fast for the first models, but as we enrich P , Δ and Γ , the number of variables generated for Sat4j and the time required for verification increases exponentially. For examples of larger sizes, it is interesting to note that Sat4j did not have enough memory to process them.

Deadlock detection as proposed in the paper is carried out in the design phase. Hence, high calculation times are not necessarily unacceptable and do not call into question the approach we propose. Let us notice that the results we present here are preliminary results and we are currently working on different ways to improve them. Specifically, we aim to integrate ELECTRUM more closely to avoid the creation of boolean variables where ELECTRUM could handle enumerated variables.

7 Conclusion and future work

We have developed a generic method to validate the proper behavior of a diagnoser at design stage. In particular, we have showed that it is possible to anticipate deadlock situations.

In this paper, we do not introduce any mechanism to identify the causes and modify the estimation model in order to eliminate deadlock scenarios, but this remains a goal for future work. In particular, there are several ways to treat a deadlock scenario. One can ignore it, as in the deadlocking path the system should stop autonomous operation and wait for operator intervention. One can alter the system's behaviour or instrumentation, which is costly and sometimes impossible. Finally, one can revise the fault management strategy, by changing the preference ordering, or simply diagnosing other elements of the system, such as whether we trust a component instead of its real health status.

In addition, we plan to reuse the automation of model-checking allowing us to verify the existence of deadlock scenarios for our diagnosers to check other properties related to diagnosis such as diagnosability.

ID	nb of robots	grid size	$\ P\ $	$\ \Delta\ $	$\ \Gamma\ $	CNF vars	CNF clauses	deadlock	no deadlock
1	1	1 x 2	13	17	7	1083	1586	0,1s	0,8s
2	1	2 x 2	21	27	11	3400	3925	0,5s	4,8s
3	1	2 x 3	29	37	15	13757	10593	7,3s	52,9s
4	2	2 x 2	34	50	14	12631	13561	6,7s	60,5s
5	2	2 x 3	46	68	18	65 030	47 669	144,6s	1001,4s

Figure 6: ELECTRUM model checking. Columns $\|P\|$, $\|\Delta\|$ and $\|\Gamma\|$ indicate respectively the number of variables of P , the number of rules of Δ and the number of preferences in Γ . Columns 7 and 8 indicate respectively the number of variables and clauses sent to solver Sat4j by ELECTRUM. The last two columns indicate the verification time for a version of the model subject to a deadlock scenario and another version for which such a scenario does not exist.

Finally, progress on model-checker performance is needed to apply it to multi-robot missions involving large models. QBF solvers [17] are potential tools for finding deadlock scenarios of bounded length, as can express the quantified form of our conditional preferences. It is thus possible to express the presence of a deadlocking path of bounded length as a QBF satisfiability problem.

References

- [1] Nuno Macedo, Julien Brunel, David Chemouil, Alcino Cunha, and Denis Kuperberg. Lightweight specification and analysis of dynamic systems with rich configurations. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering, FSE 2016*, pages 373–383, New York, NY, USA, 2016.
- [2] Meera Sampath, Raja Sengupta, Stéphane Lafortune, Kasim Sinnamohideen, and Demosthenis Teneketzis. Diagnosability of discrete-event systems. *IEEE Transactions on automatic control*, 40(9):1555–1575, 1995.
- [3] Craig Boutilier, Ronen I Brafman, Carmel Domshlak, Holger H Hoos, and David Poole. Preference-based constrained optimization with cp-nets. *Computational Intelligence*, 20(2):137–157, 2004.
- [4] Cedric Pralet, Xavier Pucel, and Stéphanie Roussel. Diagnosis of intermittent faults with conditional preferences. In *Proceedings of the 27th International Workshop on Principles of Diagnosis (DX’16)*, 2016.
- [5] Xavier Pucel and Stéphanie Roussel. Intermittent fault diagnosis as discrete signal estimation: Trackability analysis. In *Proceedings of the 28th International Workshop on Principles of Diagnosis (DX’17)*, 2017.
- [6] James Kurien and P Pandurang Nayak. Back to the future for consistency-based trajectory tracking. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 370–377, 2000.
- [7] Marie-Odile Cordier, Philippe Dague, François Lévy, Jacky Montmain, Marcel Staroswiecki, and Louise Travé-Massuyès. Conflicts versus analytical redundancy relations: a comparative analysis of the model based diagnosis approach from the artificial intelligence and automatic control perspectives. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 34(5):2163–2177, 2004.
- [8] Alexander Felfernig and Monika Schubert. Fastdiag: A diagnosis algorithm for inconsistent constraint sets. In *Proceedings of the 21st International Workshop on the Principles of Diagnosis (DX 2010)*, Portland, OR, USA, pages 31–38, 2010.
- [9] Johan De Kleer. Diagnosing multiple persistent and intermittent faults. In *Proceedings of the 21st international joint conference on Artificial intelligence*, pages 733–738. Morgan Kaufmann Publishers Inc., 2009.
- [10] Olivier Contant, Stéphane Lafortune, and Demosthenis Teneketzis. Diagnosis of intermittent faults. *Discrete Event Dynamic Systems*, 14(2):171–202, 2004.
- [11] Edmund M. Clarke, Jr., Orna Grumberg, and Doron A. Peled. *Model Checking*. MIT Press, Cambridge, MA, USA, 1999.
- [12] Craig Boutilier, Ronen I Brafman, Carmel Domshlak, Holger H Hoos, and David Poole. Cp-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res.(JAIR)*, 21:135–191, 2004.
- [13] Alessandro Cimatti, Charles Pecheur, and Roberto Cavada. Formal verification of diagnosability via symbolic model checking. In *Proceedings of the 18th international joint conference on Artificial intelligence*, pages 363–369. Morgan Kaufmann Publishers Inc., 2003.
- [14] Alessandro Cimatti, Edmund Clarke, Fausto Giunchiglia, and Marco Roveri. Nusmv: a new symbolic model checker. *International Journal on Software Tools for Technology Transfer*, 2(4):410–425, Mar 2000.
- [15] Daniel Jackson. *Software Abstractions: Logic, Language, and Analysis*. The MIT Press, 2006.
- [16] Daniel Le Berre and Anne Parrain. The sat4j library, release 2.2, system description. *Journal on Satisfiability, Boolean Modeling and Computation*, 7:59–64, 2010.
- [17] Massimo Narizzano, Luca Pulina, and Armando Tacchella. The qbfeval web portal. In *Logics in Artificial Intelligence*, pages 494–497. Springer Berlin Heidelberg, 2006.