Uncertain Temporal Observations in Diagnosis

Gianfranco Lamperti and **Marina Zanella** [1]

Abstract. In diagnosis, the notion of observation varies according to the class of considered systems. In discrete-event systems, an observation usually consists of a sequence, or a set of sequences, of totally ordered observable events. This paper extends the concept of discrete-event observation in several ways. First, observable events (messages) may be uncertain in nature, both in behavioral models and in system observations. Uncertain messages are specified by variables ranging on finite sets of observable labels. Second, messages relevant to a system observation are accommodated within a DAG, the observation graph, whose edges define a partial temporal ordering among (uncertain) messages. This way, an observation graph implicitly defines a finite set of system observations in the traditional sense. Consequently, solving a diagnostic problem amounts to solving at one time several traditional diagnostic problems. Finally, the (possibly distributed) reconstruction of the system behavior is further complicated by the fact that homonymous observable labels can be shared by different components. This raises the need of dealing with null messages. The method is appropriate for several real systems, where messages may get lost, are noisy, or attached timestamps are generated by different clocks.

1 INTRODUCTION

Model-based diagnosis (MBD) is a problem solving task aimed at troubleshooting a physical system, given an observation of the system itself and the models describing its structure and behavior. A solution produced by a diagnostic session is a set of *candidate diagnoses*, each diagnosis being a set of faulty components (or a set of specific faults assigned to components) that *explains* the observation. Thus, the system observation is an integral part of any diagnostic problem and the input of any diagnostic process. Since the beginning of MBD research it was quite clear that the temporal dimension of observation, and consequently of behavioral models, is very important for MBD [6]. However, taking the temporal dimension into account makes diagnosis significantly complex, both from the conceptual and practical point of view [3]. In temporal MBD, observation is endowed with both a logical content, expressing *what* pieces of information have been observed, and a temporal content, expressing *when* they have been observed. Dually, the notion of candidate diagnosis is twofold: it encompasses both the set of faults explaining the logical content of the observation, and the time constraints explaining the temporal location of the observation. Moreover, the logical and temporal aspects of diagnoses are closely related and cannot be generated separately. This intuition is substantiated by the notion of *explanatory diagnosis* [10].

In order to cope with the conceptual and computational difficulties of temporal MBD, several simplifying assumptions have been made in different approaches. This paper is focused on an approach that models the dynamic and/or time-varying behavior of physical systems by means of discrete state changes. A discrete-change abstraction is simpler than a continuous-change one but, at the same time, is quite powerful since, for diagnostic purposes, many continuous-variable systems can be modeled as discrete [9, 5]. In the last few years, a good deal of research efforts have been devoted to MBD of systems modeled as discrete-event [4, 12, 11, 7, 1, 2, 8]. All these approaches feature compositional modeling and analogous modeling primitives, since the behavior of each component is represented as a finite automaton.

This paper extends the diagnostic method presented in [1, 2, 8], in order to support the notion of *uncertain observation*, according to which observations may be either logically uncertain or temporally uncertain or both. An observation is *logically uncertain* if its logical content is not univocal, that is, it does not identify one set of observed labels but several ones. This includes the case when the sender of an observed label is uncertain, as homonymous observable labels may be generated by several components. An observation is *temporally uncertain* if its temporal content identifies several temporal locations of the logical content.

The solution of a diagnostic problem featuring an uncertain observation is, in principle, the union of the solutions of all the diagnostic problems inherent to the univocal observations it represents. However, the proposed method can cope with an uncertain observation uniformly, without generating and processing the single observation instances. An uncertain observation is described by a DAG, where each node is a (possibly logically uncertain) observed event and each edge is a temporal precedence relation. A logically uncertain event may be one out of a set of univocal events, where such a set may even include the null event. In previous work the observation of the system was only temporally uncertain, being a set of totally temporally ordered sequences of messages, each relevant to a single component [1, 2] and to an arbitrary group of components [8]. In the diagnoser approach [12] where, like in ours, the only temporal information that can be modeled by the available ontology of time are ordering constraints, the observation is not uncertain at all, being just one sequence of messages, pertaining to the whole system.

2 APPROACH

A discrete-event system Θ is modeled as a network of communicating finite automata, which are connected to one another through communication *links*. Automata embodies both *nominal* and *faulty* behaviors. Initially, Θ is in a *quiescent* state Θ_0. On the arrival of an event from the external world, Θ becomes *reacting*, that is, it makes a series of transitions until a final quiescent state Θ_f is reached. This *reaction* yields a number of observable events, the *messages*, which make up a system *observation* $OBS(\Theta)$. Based on $OBS(\Theta)$ and Θ_0, a *reconstruction* of the system behavior is carried out, which yields an *active space*, that is, a graph representing the whole set of histories (sequence of transitions) which explain $OBS(\Theta)$. Diagnoses are eventually generated from the active space, as described in [2].

[1] Dipartimento di Elettronica per l'Automazione, Università di Brescia, via Branze 38, 25123 Brescia, Italy, email: {lamperti,zanella}@ing.unibs.it

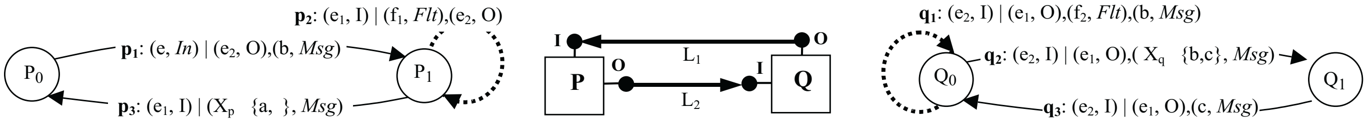
p₂: (e₁, I) | (f₁, Flt),(e₂, O)
p₁: (e, In) | (e₂, O),(b, Msg)
p₃: (e₁, I) | (Xp {a, }, Msg)
P0
P1
I
O
L1
O
I
L2
P
Q
q₁: (e₂, I) | (e₁, O),(f₂, Flt),(b, Msg)
q₂: (e₂, I) | (e₁, O),(Xq {b,c}, Msg)
q₃: (e₂, I) | (e₁, O),(c, Msg)
Q0
Q1

Figure 1. System Θ (center) and models of components P (left) and Q (right).

2.1 Component

A *component model* M_C is a communicating automaton, $M_C = (\mathcal{S}, \mathcal{E}_{in}, \mathcal{I}, \mathcal{E}_{out}, \mathcal{O}, \mathcal{T})$, where $\mathcal{S}$ is the set of states, $\mathcal{E}_{in}$ the set of inputs, $\mathcal{I}$ the set of input terminals, $\mathcal{E}_{out}$ the set of outputs, $\mathcal{O}$ the set of output terminals, and $\mathcal{T}$ the (nondeterministic) transition function, $\mathcal{T} : \mathcal{S} \times \mathcal{E}_{in} \times \mathcal{I} \times 2^{\mathcal{E}_{out} \times \mathcal{O}} \mapsto 2^{\mathcal{S}}$. A *component* is the instantiation of a component model. Each component is implicitly endowed with three *virtual* terminals, namely $In \in \mathcal{I}$, the standard input, $Msg \in \mathcal{O}$, the message terminal, and $Flt \in \mathcal{O}$, the fault terminal.

An input (output) event is a pair (E, ϑ), where E is an input (output) and ϑ an input (output) terminal. A transition T from state S_1 to state S_2 is triggered by an input event $\alpha = (E, I)$ at an input terminal I and generates the (possibly empty) set of output events $\beta = \{(E_1, O_1), \ldots, (E_n, O_n)\}$ at output terminals $O_1 \cdots O_n$, respectively. This is denoted by $T = S_1 \xrightarrow{\alpha \mid \beta} S_2$. A transition can generate at most one output on each output terminal. An event (E, Flt) is a *faulty* event and, the relevant transition, a faulty transition.

Let Λ be a set of *observable labels* and $\mathcal{V}$ a set of *variables* such that $\forall X \in \mathcal{V}$ the value of X is defined on $\|X\| \subseteq (\Lambda \cup \{\epsilon\})$, where ϵ denotes the *null* message. The set $\mu = \Lambda \cup \mathcal{V}$ is the *domain* of messages. The *extension* of a message $m \in \mu$ is the set defined as follows:

$$\|m\| = \begin{cases} \{\ell\} & \text{if } m = \ell, \ell \in \Lambda \\ \|X\| & \text{if } m = X, X \in \mathcal{V} \end{cases}$$

2.2 Link

A *link model* M_L is a 4-tuple, $M_L = (I, O, \chi, \pi)$, where I is the input terminal, O the output terminal, χ the *capacity*, $\chi \geq 1$, and π the *saturation policy*. A *link* is an instantiation of a link model, that is, the directed communication channel between two different components C and C', where an output terminal of C and an input terminal of C' coincide with the input and output terminal of L, respectively. The *state* of a link L is the queue of events in L. In fact, if a transition of C generates an event on the input terminal of L, such an event is buffered within the link. Dually, when a transition of C' is triggered by the first event in the queue within the link, such an event is dequeued. When the number of events in L equals χ, L is *saturated*. When L is saturated, the semantics for the triggering of a transition T of a component C that generates a new output event (E, O) is dictated by the saturation policy π of L, which can be either (i) *LOSE*: E is lost, (ii) *OVERRIDE*: E overrides the last event in the queue of dangling events of L, or (iii) *WAIT*: the transition T cannot be triggered until L becomes unsaturated, that is, until at least one event in L is consumed. If (E, ϑ) is an event, $Link(\vartheta)$ denotes the link relevant to ϑ.

2.3 Cluster

A *cluster* $\xi = (\mathcal{C}, \mathcal{L})$ is a connected graph where nodes are terminals of components in $\mathcal{C}$ and edges are the elements in $\mathcal{L}$, that is, the whole set of links among such terminals. A *decomposition* $\Xi = \{\xi_1, \ldots, \xi_n\}$ of ξ is a set of disjoint clusters $\xi_i = (\mathcal{C}_i, \mathcal{L}_i)$ where $\{\mathcal{C}_1, \ldots, \mathcal{C}_n\}$ is a partition of $\mathcal{C}$. The *interface* of Ξ, $Interf(\Xi)$, is the set of links $\mathcal{L}' \subseteq \mathcal{L}$ where $\forall L = C_1.O \Rightarrow C_2.I \in \mathcal{L}'$ $(C_1 \in \xi_i, C_2 \in \xi_j, i \neq j, \xi_i \in \Xi, \xi_j \in \Xi)$.

Example 1 In Figure 1, a system Θ is displayed, where P and Q are components, while L_1 and L_2 are links, for which we assume $\chi = 2$ and $\pi = LOSE$. The automata corresponding to the behavior of P and Q are displayed on the left and on the right, respectively. Both automata are composed of two states and three transitions, one of which is faulty (p_2 and q_1, respectively). For instance, transition p_1 is triggered by the input event $\alpha = (e, In)$ and generates the set of output events $\beta = \{(e_2, O), (b, Msg)\}$, where the former is directed toward Q by means of link L_2, while the latter is a message labeled b. Notice that transitions p_3 and q_2 involve uncertain messages, namely, $X_p \in \{a, \epsilon\}$ and $X_q \in \{b, c\}$. Variable X_p can be physically interpreted as an observable label a which may be lost, while the observable label associated with X_q can be either b or c. In the following, we assume Θ decomposed in $\Xi = \{\xi_p, \xi_q\}$, where $\xi_p = (\{P\}, \emptyset)$, $\xi_q = (\{Q\}, \emptyset)$, and $Interf(\Xi) = \{L_1, L_2\}$.

2.4 Observation

An *observation* $OBS(\xi)$ of a cluster $\xi = (\mathcal{C}, \mathcal{L})$ is a directed (not necessarily connected) acyclic graph, $OBS(\xi) = (\Omega, \Upsilon, \Omega_0, \Omega_f)$, called *observation graph*, such that Ω is the set of nodes, where each $\omega \in \Omega$ is marked with a message $Msg(\omega) \in \mu$, $\Upsilon : \Omega \mapsto 2^{\Omega}$ is the set of edges, $\Omega_0 \subseteq \Omega$ the set of *roots*, and $\Omega_f \subseteq \Omega$ the set of *leaves*. The '$\prec$' *temporal precedence* relationship among nodes is defined as follows:

1. If $\omega \mapsto \omega' \in \Upsilon$ then $\omega \prec \omega'$;
2. If $\omega \prec \omega'$ and $\omega' \prec \omega''$ then $\omega \prec \omega''$;
3. If $\omega \mapsto \omega' \in \Upsilon$ then $\nexists \omega'' \in \Omega \mid \omega \prec \omega'' \prec \omega'$.

Furthermore, by definition,

4. $\omega \preceq \omega'$ iff $\omega \prec \omega'$ or $\omega = \omega'$;
5. $\forall \omega_0 \in \Omega_0$ $(\nexists \omega \in \Omega \mid \omega \prec \omega_0)$;
6. $\forall \omega_f \in \Omega_f$ $(\nexists \omega \in \Omega \mid \omega_f \prec \omega)$.

An *index* $\Im$ for $OBS(\xi)$ is a subset of Ω such that $\forall \omega \in \Im$ $(\nexists \omega' \in \Im \mid \omega' \prec \omega)$. The following two functions are defined on $\Im$ (where $Cons$ stands for *consumed*):

1. $Cons(\Im) = \{\omega \mid \omega \in \Omega, \omega' \in \Im, \omega \preceq \omega'\}$;
2. $Next(\Im) = \mathcal{N} \cup \mathcal{N}^+$, where, if $\Im \neq \emptyset$, then $\mathcal{N} = \{\omega \mid \omega \in \Omega, \omega \notin Cons(\Im), \forall \omega' \mapsto \omega \in \Upsilon\ (\omega' \in Cons(\Im))\}$, else $N = \Omega_0$, and $\mathcal{N}^+ = \{\omega \mid \omega \in \Omega, \forall \omega' \in \mathcal{N}, \forall \omega'' \in \Omega, \omega' \preceq \omega'' \prec \omega\ (\epsilon \in \|Msg(\omega'')\|)\}$.

$\Im$ is *complete* when either $Cons(\Im) = \Omega$ or $\forall \omega \in (\Omega - Cons(\Im))(\epsilon \in \|Msg(\omega)\|)$.

The *restriction* of $OBS(\xi)$ on a set of components $\mathcal{C}' \subset \mathcal{C}(\xi)$, denoted by $OBS_{\langle \mathcal{C}' \rangle}(\xi)$, is an observation $(\Omega', \Upsilon', \Omega'_0, \Omega'_f)$ where, denoting with $\mu(\mathcal{C})$ the set of observable labels in Λ relevant to components in $\mathcal{C}$,

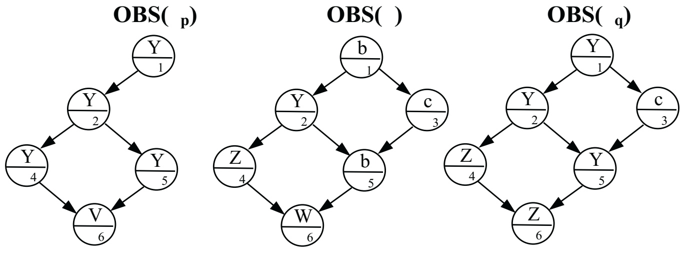

Figure 2. Observation graph of system Θ (center), and observation restrictions on clusters ξ_p (left) and ξ_q (right), where variable $Y \in \{b, \epsilon\}$, $Z \in \{b, c, \epsilon\}$, $W \in \{a, b, c\}$, and $V \in \{a, b, \epsilon\}$.

1. $\Omega' = \{\omega' \mid \omega' = \omega_{\langle \mathcal{C}' \rangle}, \omega \in \Omega\}$, where $\omega_{\langle \mathcal{C}' \rangle}$, the *restriction* of ω on $\mathcal{C}'$, is such that:
   (a) $\|Msg(\omega')\| = \{m \mid m \in \|Msg(\omega)\|\} \cup M_\epsilon(\omega)$, where
   $$M_\epsilon(\omega) = \begin{cases} \{\epsilon\} & \text{if } \exists m \ (m \in \mu(\mathcal{C}'), m \in \mu(\mathcal{C} - \mathcal{C}') \\ \emptyset & \text{otherwise} \end{cases}$$
   (b) $\|Msg(\omega')\| \neq \emptyset$;
2. $\omega'_1 \mapsto \omega'_2 \in \Upsilon', \omega'_1 \in \Omega', \omega'_2 \in \Omega', \omega'_1 = \omega_{1\langle \mathcal{C}' \rangle}, \omega'_2 = \omega_{2\langle \mathcal{C}' \rangle}$ iff
   (a) $\omega_1 \prec \omega_2$ in $OBS(\xi)$,
   (b) $\not\exists \omega'_3 \in \Omega', \omega'_3 = \omega_{3\langle \mathcal{C}' \rangle}$, such that $\omega_1 \prec \omega_3 \prec \omega_2$ in $OBS(\xi)$.

Let $\xi' = (\mathcal{C}', \mathcal{L}')$, $\mathcal{C}' \subset \mathcal{C}$. By definition, $OBS_{\langle \xi' \rangle}(\xi) \equiv OBS_{\langle \mathcal{C}' \rangle}(\xi)$. Let $\mathcal{Q}_\omega = \langle \omega_1, \ldots, \omega_p \rangle$ such that:

1. $\{\omega_1, \ldots, \omega_p\} = \Omega(OBS(\xi))$;
2. $\forall i \in [1 \,..\, p], \forall j \in [1 \,..\, p], i \neq j$ (if $\omega_i \prec \omega_j$ in $OBS(\xi)$ then ω_i precedes ω_j in $\mathcal{Q}_\omega$).

Let $\mathcal{Q}_\ell = \langle \ell_1, \ldots, \ell_p \rangle$, where $\forall i \in [1 \,..\, p](\ell_i \in \|Msg(\omega_i)\|, \omega_i \in \mathcal{Q}_\omega)$. Let $\mathcal{Q}'_\ell = \langle \ell'_1, \ldots, \ell'_{p'} \rangle, p' \leq p$, be the sequence obtained from $\mathcal{Q}_\ell$ by removing the null messages. $\mathcal{Q}'_\ell$ is an *instance* of $OBS(\xi)$. The *extension* of $OBS(\xi)$, $\|OBS(\xi)\|$, is the whole set of the instances of $OBS(\xi)$. The i-th instance of an observation can be thought of as a special case of observation, whose graph, denoted by $OBS_{(i)}(\xi)$, is a sequence of nodes corresponding to the messages in the instance.

$OBS(\xi)$ is *weaker* than $OBS'(\xi)$, namely $OBS(\xi) \sqsupset OBS'(\xi)$, iff $\|OBS(\xi)\| \supset \|OBS'(\xi)\|$. More generally, $OBS(\xi) \sqsupseteq OBS'(\xi)$ iff $\|OBS(\xi)\| \supseteq \|OBS'(\xi)\|$.

Example 2 On the center of Figure 2, an observation graph of Θ is depicted, where $\Omega = \{\omega_1, \ldots, \omega_6\}$, $\Omega_0 = \{\omega_1\}$, $\Omega_f = \{\omega_6\}$, and Υ is represented by edges among nodes. Notice that $Msg(\omega_1) = c$, while $Msg(\omega_4) = Z$, where $\|Z\| = \{b, c, \epsilon\}$. Since $\omega_1 \mapsto \omega_2 \in \Upsilon$, we have $\omega_1 \prec \omega_2$ and, since $\omega_2 \prec \omega_4$, we also have $\omega_1 \prec \omega_4$. Assuming $\Im = \{\omega_2, \omega_3\}$, we have $Cons(\Im) = \{\omega_1, \omega_2, \omega_3\}$ and $Next(\Im) = \{\omega_4, \omega_5\}$. The only complete index is $\{\omega_6\}$. The restrictions $OBS(\xi_p) = OBS_{\langle \xi_p \rangle}(\Theta)$ and $OBS(\xi_q) = OBS_{\langle \xi_q \rangle}(\Theta)$ are given on the left and on the right of Figure 2, respectively. We have $\mu(P) = \{a, b\}$ and $\mu(Q) = \{b, c\}$. Thus, P and Q share the observable label b. This is why $\omega_{1\langle \xi_p \rangle} = Y, Y \in \{b, \epsilon\}$. In fact, one cannot know a priori whether b has been actually generated by P or by Q. Since $Msg(\omega_6) = W, W \in \{a, b, c\}$, the restriction $\omega_{6\langle \xi_p \rangle}$ will include a, but not c. Besides, the ambiguous label b will be replaced by the uncertain pair (b, ϵ), so that $\omega_{6\langle \xi_p \rangle} = V$, $V \in \{a, b, \epsilon\}$. An instance of $OBS(\Theta)$ is $\langle b, c, b, c, a \rangle$, corresponding to $\mathcal{Q}_\omega = \langle \omega_1, \omega_3, \omega_2, \omega_5, \omega_4, \omega_6 \rangle$ and $\mathcal{Q}_l = \langle b, c, \epsilon, b, c, a \rangle$. It is easy to show that $\|OBS(\Theta)\|$ includes 57 instances.

2.5 Reconstruction

A pair $\wp(\xi) = (OBS(\xi), \xi_0)$, where $OBS(\xi)$ is an observation and ξ_0 the initial state of ξ, is a *diagnostic problem* for ξ. An *active space* of $\wp(\xi)$ is a finite automaton, $Act(\wp(\xi)) = (\mathcal{S}, \mathcal{E}, \mathcal{T}, S_0, \mathcal{S}_f)$, where $\mathcal{S}$ is the set of states, $\mathcal{E}$ the set of events, $\mathcal{T}$ the transition function, S_0 the initial state, and $\mathcal{S}_f \subseteq \mathcal{S}$ the set of final states, defined as follows:

1. (*Atomic active space*) if ξ incorporates a single component C with model $M_C = (\mathcal{S}_C, \mathcal{E}_{in\,C}, \mathcal{I}_C, \mathcal{E}_{out\,C}, \mathcal{O}_C, \mathcal{T}_C)$ and $OBS(\xi) = ukn$, then $\mathcal{S} = \mathcal{S}_C$, $\mathcal{E} = \mathcal{T}_C$, $\mathcal{T} : \mathcal{S} \times \mathcal{E} \mapsto \mathcal{S}$, $S_0 = \xi_0$, and $\mathcal{S}_f$ is the set of states reachable from S_0;
2. (*Compound active space*) if $\Xi = \{\xi_1, \ldots, \xi_n\}$ is a decomposition of ξ and $\mathcal{A} = \{Act(\wp(\xi_1), \ldots, Act(\wp(\xi_n)\}$ a set of active spaces relevant to Ξ, where
   (a) $\forall i \in [1 \,..\, n]$ $(Act(\wp(\xi_i) = (\mathcal{S}_i, \mathcal{E}_i, \mathcal{T}_i, S_{0_i}, \mathcal{S}_{f_i}), \wp(\xi_i) = (OBS(\xi_i), \xi_{i_0}))$, and
   (b) $OBS(\xi) \trianglerighteq \{OBS(\xi_1), \ldots, OBS(\xi_n)\}$,

   then $Act(\wp(\xi)) = \Re_{OBS(\xi)}(\mathcal{A})$, where $\Re$ is the *reconstruction operator* introduced below.

Let $\Sigma = \mathcal{S}_1 \times, \ldots, \times \mathcal{S}_n$, $\Im^*$ the domain of the indexes $\Im$ of $OBS(\xi)$, and $\mathcal{D}$ the domain of the states of the links in $Interf(\Xi)$. The *spurious active space* of $\mathcal{A}$ based on the observation $OBS(\xi)$, obtained by applying the *spurious reconstruction operator* $\widetilde{\Re}$, is an automaton $\widetilde{\Re}_{OBS(\xi)}(\mathcal{A}) = (\widetilde{\mathcal{S}}, \mathcal{E}, \widetilde{\mathcal{T}}, S_0, \mathcal{S}_f)$, where

- $\widetilde{\mathcal{S}} = \{S_0\} \cup \{S' \mid S \xrightarrow{\mathcal{T}} S' \in \widetilde{\mathcal{T}}\}, \widetilde{\mathcal{S}} \subseteq \Sigma \times \Im^* \times \mathcal{D}$;
- $\mathcal{E} = \bigcup_{i=1}^{n} \mathcal{E}_i$;
- $\widetilde{\mathcal{T}} : \widetilde{\mathcal{S}} \times \mathcal{E} \mapsto \widetilde{\mathcal{S}}$;
- $S_0 = (\sigma_0, \Im_0, D_0)$, where $\sigma_0 = (S_{0_1}, \ldots, S_{0_n})$, $\Im_0 = \emptyset, D_0 = (\emptyset, \ldots, \emptyset)$;
- $\mathcal{S}_f = \{(\sigma_f, \Im_f, D_f)$, where $\forall S_i \in \sigma_f \ (S_i \in \mathcal{S}_{f_i}), \Im_f$ is complete, and $D_f = (\emptyset, \ldots, \emptyset)\}$.

The transition function $\widetilde{\mathcal{T}}$ is defined as follows: $N \xrightarrow{T} N' \in \widetilde{\mathcal{T}}$, where $N = (\sigma, \Im, D)$, $T = S \xrightarrow{\alpha|\beta} S'$, $N' = (\sigma', \Im', D')$, $\sigma = (S_1, \ldots, S_n), \sigma' = (S'_1, \ldots, S'_n)$ iff:

1. T is *triggerable*, that is, defining
   - $L_\alpha = Link(\theta_\alpha), \alpha = (E_\alpha, \theta_\alpha), L_\alpha \in Interf(\Xi)$,
   - $\mathcal{L}_\beta = \{L_\beta = Link(B) \mid (E_B, B) \in \beta, L_\beta \in Interf(\Xi)\}$,
   - $\mathcal{L}^u_\beta = \{L_\beta \in \mathcal{L}_\beta \mid L_\beta$ is not saturated or $L_\beta = L_\alpha\}$,
   - $\mathcal{L}^s_\beta = \mathcal{L}_\beta - \mathcal{L}^u_\beta$,
   - $\mathcal{L}^{so}_\beta = \{L_\beta \in \mathcal{L}^s_\beta \mid \pi(L_\beta) = OVERRIDE\}$,
   - $\mathcal{L}^{sw}_\beta = \{L_\beta \in \mathcal{L}^s_\beta \mid \pi(L_\beta) = WAIT\}$,

   and denoting with $L(D)[i]$ the i-th event within link L of D, we have:
   (a) $L_\alpha(D)[1] = E_\alpha$,
   (b) One of the following conditions holds:
      i. $(m, Msg) \notin \beta$
      ii. $(m, Msg) \in \beta$ and $\epsilon \in \|m\|$
      iii. $(m, Msg) \in \beta$ and $(\|m\| \cap \{\ell \mid \ell \in \|Msg(\omega)\|, \omega \in Next(\Im)\}) \neq \emptyset$,
   (c) $\mathcal{L}^{sw}_\beta = \emptyset$;
2. σ' is such that $\forall i \in [1 \,..\, n]$ (if $T \in \mathcal{T}_i$ then $S'_i = S'$ else $S'_i = S_i$);
3. If $(m, Msg) \notin \beta$ then $\Im' = \Im$ else $\Im' \in \Im^+$, where $\Im^+$ is the smallest set of indexes defined by the following rules:

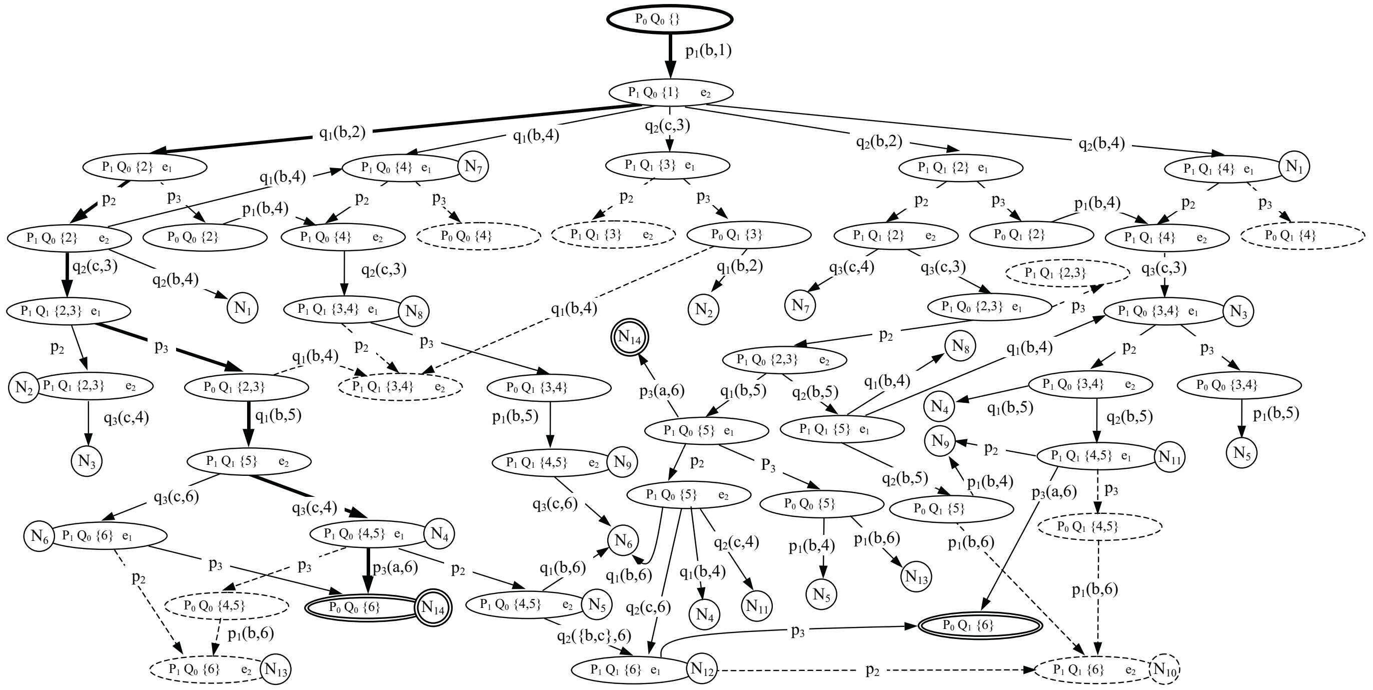

Figure 3. Active space corresponding to the canonical reconstruction $\widehat{\Re}(OBS(\Theta), \Theta_0)$, where $\Theta_0 = (P_0, Q_0)$.

(a) If $\epsilon \in \|m\|$ then $\Im \in \Im^+$,

(b) $\Im_m \in \Im^+$, where

- $\Im_m = (\Im \cup \{\omega_m\}) - \{\omega \mid \omega \in \Im, \omega \prec \omega_m\}$,
- $\omega_m \in \{\omega \mid \omega \in Next(\Im), ((\|m\| - \{\epsilon\}) \cap \|Msg(\omega)\|) \neq \emptyset\}$;

4. D' is such that, denoting with $|L(D)|$ the number of events within link L of D, the following conditions hold:

(a) $|L_\alpha(D')| = |L_\alpha(D)| - 1, \forall x \in [1 .. (|L_\alpha(D)| - 1)]$ $(L_\alpha(D')[x] = L_\alpha(D)[x+1])$;

(b) $\forall (E, \vartheta) \in \beta, L_\beta = Link(\vartheta), L_\beta \in \mathcal{L}^u_\beta$ $(L_\beta(D')[|L_\beta(D)| + 1] = E, |L_\beta(D')| = |L_\beta(D)| + 1)$;

(c) $\forall (E, \vartheta) \in \beta, L_\beta = Link(\vartheta), L_\beta \in \mathcal{L}^{so}_\beta$ $(L_\beta(D')[|L_\beta(D)|] = E, |L_\beta(D')| = |L_\beta(D)|)$;

(d) $\forall L \in (Interf(\Xi) - (\mathcal{L}_\alpha \cup \mathcal{L}^u_\beta \cup \mathcal{L}^{so}_\beta))$ $(L(D') = L(D))$.

A state $S \in \widetilde{\mathcal{S}}$ is *convergent* if there exists a path $S \rightsquigarrow S_f$ in $\widetilde{\Re}_{OBS(\xi)}(\mathcal{A})$ such that $S_f \in \mathcal{S}_f$. A transition $S \xrightarrow{\tau} S' \in \widetilde{\mathcal{T}}$ is *convergent* if S and S' are convergent. The *reconstruction* $\Re$ of $\mathcal{A}$ based on the observation $OBS(\xi)$ is the automaton $\Re_{OBS(\xi)}(\mathcal{A}) = (\mathcal{S}, \mathcal{E}, \mathcal{T}, S_0, \mathcal{S}_f)$ obtained by removing from $\widetilde{\Re}$ the states and transitions which are not convergent.

Let $\mathcal{A}$ be the set of atomic active spaces relevant to components in ξ. The *canonical reconstruction* of $\wp(\xi)$, $\widehat{\Re}(\wp(\xi))$, is the reconstruction based on the atomic active spaces, that is, $\widehat{\Re}(\wp(\xi)) = \Re_{OBS(\xi)}(\mathcal{A})$.

A path $S_0 \rightsquigarrow S_f$ in $Act(\wp(\xi))$, where $S_f \in \mathcal{S}_f$, is a *history* of ξ. The whole (possibly boundless) set of histories in $Act(\wp(\xi))$ is denoted by $\mathcal{H}(Act(\wp(\xi)))$.

Example 3 Consider the diagnostic problem $\wp(\Theta) = (OBS(\Theta), \Theta_0)$, where $OBS(\Theta)$ is depicted in Figure 2 and $\Theta_0 = (P_0, Q_0)$ (see Figure 1). The canonical reconstruction $A = \widehat{\Re}(\wp(\Theta))$, based on the atomic active spaces of components P and Q, is shown in Figure 3, where ellipses and edges represent states and transitions of A, respectively. The displayed graph corresponds to the search carried out by an extension of the *Reconstruct* algorithm presented in [1]. Dashed lines denote backtracking. We assume the decomposition $\Xi = \{\xi_p, \xi_q\}$. Each state is characterized by three fields, these being the states of P and Q (namely, σ), the index of $OBS(\Theta)$ (namely, $\Im$), and the state of the links in $Interf(\Xi) = \{L_1, L_2\}$ (namely, D). The initial state (denoted in bold) is $(\sigma_0 = (P_0, Q_0), \Im_0 = \{\}, D_0 = (\langle\rangle, \langle\rangle))$. Each edge in A is marked with the identifier of the corresponding transition in the atomic active space, possibly followed by an observable label and the identifier of a relevant node in $OBS(\Theta)$. For example, the first edge is marked by $p_1(b, 1)$, which means that, a transition p_1 of component P_1 is triggered, assuming that it produces the observable label b associated with node ω_1 in $OBS(\Theta)$. Note that, to simplify the intricacy of the graph, several nodes have been identified by labels N_i, which are then used for indirect referencing (circles), e.g., $N_{14} = ((P_0, Q_0), \{6\}, (\langle\rangle, \langle\rangle))$. Final states are denoted by double nodes. In Figure 3, the path in bold from the initial state to the final state N_{14} is one of the histories of Θ, namely $h = \langle p_1(b, \omega_1),\ q_1(b, \omega_2),\ p_2,\ q_2(c, \omega_3),\ p_3,\ q_1(b, \omega_5),\ q_3(c, \omega_4),\ p_3(a, \omega_6)\rangle$. Notice that the sequence of messages relevant to h is an instance of $OBS(\Theta)$, that is, $\langle b, b, c, b, c, a\rangle \in \|OBS(\Theta)\|$. In other words, h explains $OBS(\Theta)$. Incidentally, transition from N_5 to N_{12}, namely $q_2(\{b, c\}, 6)$, incorporates two distinct labels, b and c. This means that the set $\{b, c\}$ is shared by $\|X_q\| = \{b, c\}$, in the behavioral model of Q, and $\|W\| = \{a, b, c\}$ in ω_6. Therefore, both b and c might have been generated by transition q_1 during the reaction. According to the definition of active space, besides providing the explanation of the observation, a history is required to involve transitions which do not violate the management policy of links and, for a transition to be triggered by an internal input event, such an event is expected to be ready in the relevant link of Θ. For example, in transition $p_3(a, 6)$ from N_4 to N_{14}, event e_1 is actually within link L_1, as specified in $N_4 = ((P_1, Q_0), \{4, 5\}, (\langle e_1 \rangle, \langle\rangle))$. A node which is not within a history is inconsistent. For example, $N_{10} = ((P_1, Q_1), \{6\}, (\langle\rangle, \langle e_2 \rangle))$ is inconsistent because transition q_3, which might be activated by event e_2, would generate the message c, which is inconsistent with index $\Im = \{\omega_6\}$. In fact, $Cons(\{\omega_6\}) = \Omega$, which means that $\Im$ is complete.

Proposition 1 *Let $\wp(\xi) = (OBS(\xi), \xi_0)$ be a diagnostic problem where $\|OBS(\xi)\| = \bigcup_{i=1}^{n} \{OBS_{(i)}(\xi)\}$ is the set of instances of $OBS(\xi)$. Then,*

$$\mathcal{H}(\widehat{\Re}(\wp(\xi))) = \bigcup_{i=1}^{n} \mathcal{H}(\widehat{\Re}(OBS_{(i)}(\xi), \xi_0)).$$

Formally, Proposition 1 asserts the soundness and completeness of the diagnostic method. In other words, it claims that solving a diagnostic problem $\wp(\xi)$ involving an uncertain observation $OBS(\xi)$ produces the same histories as solving the n diagnostic problems corresponding to the whole set of totally ordered observations implicitly incorporated in $OBS(\xi)$.

Proposition 2 *Let $OBS(\xi)$ be an observation of a cluster ξ with initial state ξ_0, $\Xi = \{\xi_1, \ldots, \xi_n\}$ a decomposition of ξ, and $\hat{\mathcal{R}} = \bigcup_{i=1}^{n} \{\widehat{\Re}(OBS(\xi_i), \xi_{i_0}\}$ a set of relevant canonical reconstructions such that $\xi_0 = (\xi_{1_0}, \ldots, \xi_{n_0})$ and $\forall i \in [1 \,..\, n]$ $(OBS(\xi_i) = OBS_{\langle \xi_i \rangle}(\xi))$. Then,*

$$\mathcal{H}(\Re_{OBS(\xi)}(\hat{\mathcal{R}})) = \mathcal{H}(\widehat{\Re}(OBS(\xi), \xi_0)).$$

Proposition 2 opens the way to modular behavior reconstruction, by asserting the equivalence between a canonical active space and a compound active space built from lower-level (canonical) partial active spaces. More generally, Proposition 2, which is the theoretical basis for the (distributed and parallel) stepwise reconstruction of the active space based on a *reconstruction plan*, as introduced in [2] and formalized in [8], is still valid for discrete-event systems characterized by uncertain observations.

2.6 Diagnosis

A history h_f in an active space $Act(\wp(\xi))$ is *faulty* if and only if h_f includes at least one faulty transition T_f. The component relevant to T_f is a faulty component. A *diagnosis* δ of $\wp(\xi)$ is the set of faulty components relevant to a history in $Act(\wp(\xi))$. The set of *candidate diagnoses* of $\wp(\xi)$, namely $\Delta(\wp(\xi)) = \{\delta_1, \ldots, \delta_n\}$, is the whole set of diagnoses relevant to the histories in $Act(\wp(\xi))$.

Example 4 Considering the history $h = \langle p_1(b, \omega_1), q_1(b, \omega_2), p_2, q_2(c, \omega_3), p_3, q_1(b, \omega_5), q_3(c, \omega_4), p_3(a, \omega_6) \rangle$ highlighted in Figure 3, since p_2 and q_1 are faulty transitions for P and Q, respectively, the diagnosis relevant to h is $\delta(h) = \{P, Q\}$. Alternatively, a diagnosis can be defined in terms of faulty events relevant to a history, named *deep diagnosis* [2]. The deep diagnosis relevant to h is $\delta^d(h) = \{(P, f_1), (Q, f_2)\}$.

3 CONCLUSION

The choice of the notion of observation to adopt for temporal MBD depends on the application domain. If all the observable labels generated by a system are received by the observer, and the latter can trace the sender component of each label, then there is no need for a logically uncertain notion of observation. If, instead, a number of labels may get lost during the transmission (e.g. owing to masking phenomena), or the receiver cannot further discriminate the value of each label within a set of values (e.g. owing to noise), or homonymous labels may be generated by distinct components, then the set of labels received by the observer differs from those generated by the system, and/or labels can be ascribed to different sets of components. Therefore, the received labels cannot be described as a univocal set, but, rather, as a concise representation of all the possible sets of labels generated by the system that are compatible with it. This is actually what we call a logically uncertain notion of observation. The diagnostic task has to supply as output the candidate diagnoses for all the corresponding diagnostic problems. Analogously, if all the temporal constraints holding when the observable labels are generated, reach the observer, then a temporally uncertain notion of observation is useless. If, instead, only a subset of the temporal constraints holding when the observable labels are generated is known to the observer, then a temporally uncertain notion of observation has to be adopted. This case is typical when labels are sent to the observer by means of distinct channels and the temporal constraints holding among the labels sent on different channels get lost, or when the timestamps associated with observable labels are produced by different clocks. Since these clocks may be not perfectly synchronized, the relative order between two timestamps t and t' generated by two different clocks is uncertain when $|t - t'| < E$, where E is the time difference between the two clocks. Thus, no relative ordering relationship has to be assumed for such events.

This paper extends the modeling primitives and diagnostic method of a previous research [1, 2, 8], in order to support observations that are uncertain. The proposed notion of observation widens the class of physical systems that can be dealt with by the quoted approach, allowing for the representation of several sources of uncertainty due to partial system observability.

REFERENCES

[1] P. Baroni, G. Lamperti, P. Pogliano, and M. Zanella, ‘Diagnosis of active systems’, in *ECAI-98*, pp. 274–278, Brighton, UK, (1998).

[2] P. Baroni, G. Lamperti, P. Pogliano, and M. Zanella, ‘Diagnosis of large active systems’, *Artificial Intelligence*, **110**(1), 135–183, (1999).

[3] V. Brusoni, L. Console, P. Terenziani, and D. Theseider Dupré, ‘A spectrum of definitions for temporal model-based diagnosis’, *Artificial Intelligence*, **102**(1), 39–80, (1998).

[4] M.O. Cordier and S. Thiébaux, ‘Event-based diagnosis for evolutive systems’, in *Fifth International Workshop on Principles of Diagnosis*, pp. 64–69, New Paltz, NY, (1994).

[5] D. Förstner and J. Lunze, ‘Qualitative modelling of a power stage for diagnosis’, in *Thirteenth International Workshop on Qualitative Reasoning*, pp. 105–112, Lock Awe, UK, (1999).

[6] W. Hamscher and R. Davis, ‘Diagnosing circuits with state: an inherently underconstrained problem’, in *AAAI-84*, pp. 142–147, Austin, TX, (1984).

[7] G. Lamperti and P. Pogliano, ‘Event-based reasoning for short circuit diagnosis in power transmission networks’, in *IJCAI-97*, pp. 446–451, Nagoya, J, (1997).

[8] G. Lamperti and M. Zanella, ‘Diagnosis of discrete-event systems integrating synchronous and asynchronous behavior’, in *Tenth International Workshop on Principles of Diagnosis*, pp. 129–139, Loch Awe, UK, (1999).

[9] J. Lunze, ‘Qualitative modelling of dynamical systems: Motivation, methods, and prospective applications’, *Mathematics and Computers in Simulation*, **46**, 465–483, (1998).

[10] S.A. McIlraith, ‘Explanatory diagnosis: conjecturing actions to explain observations’, in *Sixth International Conference on Principles of Knowledge Representation and Reasoning*, pp. 167–177, Trento, I, (1998).

[11] L. Rozé, ‘Supervision of telecommunication network: a diagnoser approach’, in *Eighth International Workshop on Principles of Diagnosis*, Mont St. Michel, F, (1997).

[12] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D.C. Teneketzis, ‘Failure diagnosis using discrete-event models’, *IEEE Transactions on Control Systems Technology*, **4**(2), 105–124, (1996).